Primitive Words, Free Factors and Measure Preservation

Doron Puder*
Einstein Institute of Mathematics
Hebrew University, Jerusalem
doronpuder@gmail.com

December 21, 2012

Abstract

Let $\mathbf{F}_k$ be the free group on k generators. A word $w \in \mathbf{F}_k$ is called primitive if it belongs to some basis of $\mathbf{F}_k$. We investigate two criteria for primitivity, and consider more generally, subgroups of $\mathbf{F}_k$ which are free factors.

The first criterion is graph-theoretic and uses Stallings core graphs: given subgroups of finite rank $H \leq J \leq \mathbf{F}_k$ we present a simple procedure to determine whether H is a free factor of J. This yields, in particular, a procedure to determine whether a given element in $\mathbf{F}_k$ is primitive.

Again let $w \in \mathbf{F}_k$ and consider the word map $w : G \times \ldots \times G \to G$ (from the direct product of k copies of G to G), where G is an arbitrary finite group. We call w *measure preserving* if given uniform measure on $G \times \ldots \times G$, w induces uniform measure on G (for every finite G). This is the second criterion we investigate: it is not hard to see that primitivity implies measure preservation and it was conjectured that the two properties are equivalent. Our combinatorial approach to primitivity allows us to make progress on this problem and in particular prove the conjecture for $k = 2$.

It was asked whether the primitive elements of $\mathbf{F}_k$ form a closed set in the profinite topology of free groups. Our results provide a positive answer for $\mathbf{F}_2$.
Keywords: word maps, primitive elements of free groups, primitivity rank

1 Introduction

An element w of a free group J is called *primitive* if it belongs to some basis (free generating set) of J. When J is given with a basis X, this is equivalent to the existence of an automorphism of J which sends w to a given element of X.

*Supported by Advanced ERC Grant 247034 of Aner Shalev, and by Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

The notion of primitivity has a natural extension to subgroups in the form of free factors. Let H be a subgroup of the free group J (in particular, H is free as well). We say that H is a *free factor* of J and denote $H \overset{*}{\leq} J$, if there is another subgroup $H' \leq J$ such that $H * H' = J$. Equivalently, $H \overset{*}{\leq} J$ if every basis of H can be extended to a basis of J. (This in turn is easily seen to be equivalent to the condition that *some* basis of H extends to a basis of J).

Let $\mathbf{F}_k$ be the free group on k generators with a fixed basis $X = \{x_1, \ldots, x_k\}$. We study finitely generated subgroups of $\mathbf{F}_k$ (denoted $H \leq_{fg} \mathbf{F}_k$) and relations among them using core graphs, also known as Stallings' graphs (See [Sta83]. Actually our definition is a bit different than Stalling's, see below). Associated with every subgroup $H \leq \mathbf{F}_k$ is a directed, pointed, edge-labeled graph denoted $\Gamma_X(H)$. Edges are labeled by the elements of the given basis $X = \{x_1, \ldots, x_k\}$ of $\mathbf{F}_k$. A full definition appears in Section 2, but we illustrate the concept in Figure 1.1. It shows the core-graph of the subgroup of $\mathbf{F}_2$ generated by $x_1 x_2 x_1^{-1} x_2^{-1}$ and $x_2 x_1^2$.

Figure 1.1: The core graph $\Gamma_X(H)$ where $H = \langle x_1 x_2^{-1} x_1, x_1^{-2} x_2 \rangle \leq \mathbf{F}_2$.

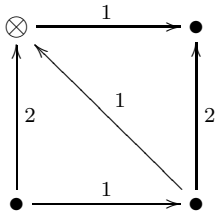

Core graphs are a key tool in the research of free groups, and are both used for proving new results and for introducing simple proofs to known results (see, for instance, [KM02, MVW07], for a survey of many such results and for further references).

A central new ingredient of our work is a new perspective on core graphs. There is a naturally defined notion of quotient on such graphs (see Section 3). In particular, we introduce in Section 3 the notion of *immediate quotients*. This in turn yields a directed graph whose vertices are all core graphs of finitely generated subgroups of $\mathbf{F}_k$ (w.r.t. the fixed basis X). A directed edge in this graph stands for the relation of an immediate quotient. This is a directed acyclic graph (DAG) i.e., it contains no directed cycles. As always, reachability in a DAG induces a distance function between vertices. Namely $\rho_X(x, y)$ is the shortest length of a directed path from x to y. We mention that the transitive closure of the immediate quotient relation is the relation "being a quotient of" which is a partial order (a lattice, in fact) on all core graphs of f.g. subgroups of $\mathbf{F}_k$. The following theorem gives a simple criterion for free factorness in terms of this distance:

Theorem 1.1. *Let $H, J \leq_{fg} \mathbf{F}_k$, and assume $\Gamma_X(J)$ is a quotient of $\Gamma_X(H)$. Then $H \overset{*}{\leq} J$ if and only if*

$$\rho_X(H, J) = rk(J) - rk(H)$$

We note that $\rho_X(\cdot,\cdot)$ can be explicitly computed, and this theorem thus yields automatically an algorithm to determine, for two given $H, J \leq_{fg} \mathbf{F}_k$ whether H is a free factor of J. In particular, it can serve to detect primitive words (see Appendix A). More generally, for any f.g. free groups $H \leq J$, this theorem can serve to detect the minimal number of complementary generators needed to obtain J from H (Corollary 3.6).

In fact, the core graph of every $H \leq_{fg} \mathbf{F}_k$ has finitely many quotients (or reachable vertices). This set is also known in the literature as the *fringe* of H (see, e.g. [MVW07]). For example, Figure 3.1 shows the fringe of the subgroup $H = \langle [x_1, x_2] \rangle$. The difference in ranks between H and $\mathbf{F}_2$ is 1. However, the distance between the corresponding core graphs in the fringe is 2. This proves that H is not a free factor of $\mathbf{F}_2$, or equivalently that $[x_1, x_2]$ is not primitive. We elaborate more in Appendix A.1.

Remark 1.2. We stress that there are other graph-theoretic algorithms to detect free factors and primitive words, including simplifications of the seminal Whitehead algorithm (the algorithm first appeared in [Whi36a, Whi36b], for its graph-theoretic simplifications see [Ger84, Sta99]). Our approach, however, is very different and does not rely on Whitehead automorphisms. We elaborate more on this in Appendix A.

Theorem 1.1 is also used for the other concept we study here, that of measure preservation of word maps. Associated with every $w \in \mathbf{F}_k$ is a *word map*. We view w as a word in the letters of the basis X. For every group G, this mapping which we also denote by w maps $\underbrace{G \times G \times \cdots \times G}_{k} \longrightarrow G$ as follows: It maps the k-tuple $(g_1, \ldots, g_k)$ to the element $w(g_1, \ldots, g_k) \in G$, where $w(g_1, \ldots, g_k)$ is the element obtained by replacing $x_1, \ldots, x_k$ with $g_1, \ldots, g_k$ (respectively) in the expression for w, and then evaluating this expression as a group element in G.

During the last years there has been a great interest in word maps in groups, and extensive research was conducted (see, for instance, [Sha09], [LS09]; for a recent book on the topic see [Seg09]). Our focus here is on the property of measure preservation: We say that the word w preserves measure with respect to a finite group G if when k-tuples of elements from G are sampled uniformly, the image of the word map w induces the uniform distribution on G. (In other words, all fibers of the word map have the same size). We say that w is *measure preserving* if it preserves measure with respect to *every* finite group G.

This concept was investigated in several recent works. See for example [LS08] and [GS09], where certain word maps are shown to be almost measure preserving, in the sense that the distribution induced by w on finite simple groups G tends to uniform, say, in L_1 distance, when $|G| \to \infty$.

Measure preservation can be equivalently defined as follows: fix some finite group G, and select a homomorphism $\alpha_G \in Hom(\mathbf{F}_k, G)$ uniformly at random. A homomorphism from a free group is uniquely determined by choosing the images of the elements of a basis, so that every homomorphism is chosen with probability $1/|G|^k$.

We then say that $w \in \mathbf{F}_k$ is measure preserving if for every finite group G and a random homomorphism α_G as above, $\alpha_G(w)$ is uniformly distributed over G.

We note that there is a stronger condition of measure preservation on a word w that is discussed in the literature. In this stronger condition we consider the image of w over the broader class of compact groups G w.r.t. their Haar measure. Our results make use only of the weaker condition that involves only finite groups.

Measure preservation can also be defined for f.g. subgroups.

Definition 1.3. For $H \leq_{fg} \mathbf{F}_k$ we say that H is *measure preserving* iff for any finite group G and $\alpha_G \in Hom(\mathbf{F}_k, G)$ a randomly chosen homomorphism as before, $\alpha_G|_H$ is uniformly distributed in $Hom(H, G)$.

In particular, $1 \neq w \in \mathbf{F}_k$ is measure preserving iff $\langle w \rangle$ is measure preserving.

It is easily seen that primitivity or free factorness yield measure preservation. The reason is that as mentioned, a homomorphism in $Hom(\mathbf{F}_k, G)$ is completely determined by the images of the elements of a basis of $\mathbf{F}_k$, which can be chosen completely arbitrarily and independently.

Several authors have conjectured that the converse is also true:

Conjecture 1.4. *For every $w \in \mathbf{F}_k$,*

$$w \text{ is primitive } \iff w \text{ is measure preserving}$$

More generally, for $H \leq_{fg} \mathbf{F}_k$,

$$H \overset{*}{\leq} \mathbf{F}_k \iff H \text{ is measure preserving}$$

From private conversations we know that this has occurred to the following mathematicians and discussed among themselves: T. Gelander, A. Shalev, M. Larsen and A. Lubotzky. The question was mentioned several times in the Einstein Institute Algebra Seminar. This conjecture was independently raised in [LP10][1].

Here we prove a partial result:

Theorem 1.5. *Let $H \leq_{fg} \mathbf{F}_k$ have rank $\geq k-1$. Then,*

$$H \overset{*}{\leq} \mathbf{F}_k \iff H \text{ is measure preserving}$$

In particular, for every $w \in \mathbf{F}_2$:

$$w \text{ is primitive } \iff w \text{ is measure preserving}$$

[1]It is interesting to note that there is an easy abelian parallel to Conjecture 1.4: A word $w \in \mathbf{F}_k$ is primitive, i.e. belongs to a basis, in $\mathbb{Z}^k \cong \mathbf{F}_k/\mathbf{F}_k'$ iff for any group G the associated word map is surjective. See [Seg09], Lemma 3.1.1.

The proof of this result relies, inter alia, on Theorem 1.1. Note that a set of $k-1$ elements $w_1, \ldots, w_{k-1} \in \mathbf{F}_k$ can be extended to a basis iff it is a free set that generates a free factor. Thus, the result for subgroups can also be stated for finite subsets as follows: Let $r \geq k-1$. A set $\{w_1, \ldots, w_r\} \subset \mathbf{F}_k$ can be extended to a basis iff for every finite group G and random homomorphism α_G as above, the r-tuple $(\alpha_G(w_1), \ldots, \alpha_G(w_r))$ is uniformly distributed in G^r, the direct product of r copies of G.

There is an interesting connection between this circle of ideas and the study of profinite groups. For example, an immediate corollary of Theorem 1.5 is that

Corollary 1.6. *The set of primitive elements in* $\mathbf{F}_2$ *is closed in the profinite topology.*

We discuss this corollary and other related results in Section 7.

In order to prove Conjecture 1.4, one needs to find for every non-primitive word $w \in \mathbf{F}_k$, some witness finite group G with respect to which w is not measure preserving. Our witnesses are always the symmetric groups S_n.

It is conceivable that our method of proof for Theorem 1.5 is powerful enough to establish Conjecture 1.4. We define two categorizations of elements (and of f.g. subgroups) of free groups $\pi(\cdot)$ and $\phi(\cdot)$. They map every free word and free subgroup into $\{0, 1, 2, 3, \ldots\} \cup \{\infty\}$. We believe these two maps are in fact identical. This, if true, yields the general conjecture. Presently we can show that they are equivalent under certain conditions, and this yields our partial result.

The first categorization is called *the primitivity rank*. It is a simple fact that if $w \in \mathbf{F}_k$ is primitive, then it is also primitive in every subgroup of $\mathbf{F}_k$ containing it (see Claim 2.5). However, if w is not primitive in $\mathbf{F}_k$, it may be either primitive or non-primitive in subgroups containing it. But what is the smallest rank of a subgroup in which we can realize w is not primitive? Informally, how far does one have to search in order to establish that w is *not* primitive in $\mathbf{F}_k$? Concretely:

Definition 1.7. The **primitivity rank** of $w \in \mathbf{F}_k$, denoted $\pi(w)$, is

$$\pi(w) = min \left\{ rk(J) \;\middle|\; \begin{array}{c} w \in J \leq \mathbf{F}_k \; s.t. \\ w \text{ is } \mathbf{not} \text{ primitive in } J. \end{array} \right\}$$

If no such J exists, $\pi(w) = \infty$. A subgroup J for which the minimum is obtained is called w-**critical**.

This extends naturally to subgroups. Namely,

Definition 1.8. For $H \leq_{fg} \mathbf{F}_k$, the primitivity rank of H is

$$\pi(H) = min \left\{ rk(J) \;\middle|\; \begin{array}{c} H \leq J \leq \mathbf{F}_k \; s.t. \\ H \text{ is } \mathbf{not} \text{ a free factor of } J. \end{array} \right\}$$

Again, if no such J exists, $\pi(H) = \infty$. A subgroup J for which the minimum is obtained is called H**-critical**.

For instance, $\pi(w) = 1$ if and only if w is a proper power of another word (i.e. $w = v^d$ for some $v \in \mathbf{F}_k$ and $d \geq 2$). In Section 4 we show (Corollary 4.2) that in $\mathbf{F}_k$ the primitivity rank takes values only in $\{0, 1, 2, \ldots, k\} \cup \{\infty\}$ (the only word w with $\pi(w) = 0$ is $w = 1$). Lemma 4.1 shows, moreover, that $\pi(w) = \infty$ ($\pi(H) = \infty$, resp.) iff w is primitive ($H \overset{*}{\leq} \mathbf{F}_k$). Finally Lemma 6.8 yields that π can take on every value in $\{0, \ldots, k\}$. For example, if $\mathbf{F}_k$ is given with some basis $X = \{x_1, \ldots, x_k\}$ then for every $1 \leq d \leq k$, $\pi(x_1^2 \ldots x_d^2) = d$. It is interesting to mention that $\pi(H)$ also generalizes the notion of *compressed* subgroups, as appears, e.g., in [MVW07]: a subgroup $H \leq_{fg} \mathbf{F}_k$ is compressed iff $\pi(H) \geq rk(H)$.

The second categorization of sets of formal words has its roots in [Nic94] and more explicitly in [LP10]. It concerns homomorphisms from $\mathbf{F}_k$ to the symmetric groups S_n, and more concretely the probability that 1 is a fixed point of the permutation $w(\sigma_1, \ldots, \sigma_k)$ for some $w \in \mathbf{F}_k$ when $\sigma_1, \ldots, \sigma_k \in S_n$ are chosen randomly with uniform distribution. More generally, for a subgroup $H \leq_{fg} \mathbf{F}_k$ we study the probability that 1 is a common fixed point of (the permutations corresponding to) all elements in H. We ask how much this probability deviates from the corresponding probability in the case of measure preserving subgroups, i.e. from $\frac{1}{n^{rk(H)}}$. (We continue the presentation for subgroups only. This clearly generalizes the case of a word: for every word $w \neq 1$ consider the subgroup $\langle w \rangle$.)

Formally, for $H \leq_{fg} \mathbf{F}_k$ we define the following function whose domain is all integers $n \geq 1$ where $\alpha_n \in Hom(\mathbf{F}_k, S_n)$ is a random homomorphism with uniform distribution:

$$\Phi_H(n) = Prob\big[\forall w \in H \quad \alpha_n(w)(1) = 1\big] - \frac{1}{n^{rk(H)}} \tag{1.1}$$

Clearly, if H is measure preserving, then Φ_H vanishes for every $n \geq 1$.

Nica [Nic94] showed that for a fixed word $w \neq 1$ and large enough n, it is possible to express $\Phi_w(n)$ ($=\Phi_{\langle w \rangle}(n)$) as a rational function in n. We show below that this is easily extended to apply to $\Phi_H(n)$ for arbitrary $H \leq_{fg} \mathbf{F}_k$. Nica's clever observation was used in [LP10] to introduce a new categorization of free words, denoted $\phi(\cdot)$, which, like $\pi(\cdot)$, associates a non-negative integer or ∞ to every formal word (note that in [LP10] the notion of primitive words has a different meaning than in the current paper). This categorization can also be extended to arbitrary finitely generated subgroups of $\mathbf{F}_k$. More specifically, it is shown in Section 5 that for every $H \leq_{fg} \mathbf{F}_k$ and n large enough (say, at least the number of vertices in the core graph of H), we have

$$\Phi_H(n) = \sum_{i=0}^{\infty} a_i(H) \frac{1}{n^i} \tag{1.2}$$

where the coefficients $a_i(H)$ are integers depending only on H. We define $\phi(H)$ as follows:

$$\phi(H) := \begin{cases} \text{the smallest integer } i \text{with } a_i(H) \neq 0 & \text{if } \Phi_H(n) \not\equiv 0 \\ \infty & \text{if } \Phi_H(n) \equiv 0 \end{cases} \tag{1.3}$$

Thus, $\phi(H)$ measures to what extent the probability that 1 is a common fixed point of H differs from $\frac{1}{n^{rk(H)}}$, the corresponding probability if H were measure preserving. The higher $\phi(H)$ is, the closer the probability is asymptotically to $\frac{1}{n^{rk(H)}}$. If H is a measure preserving subgroup, then $\phi(H) = \infty$.

As it turns out there is a strong connection between $\pi(H)$ and $\phi(H)$. Already Nica's result can be interpreted in the language of $\phi(\cdot)$ to say that $\phi(w) = 1$ iff w is a power, that is iff $\pi(w) = 1$. But the connection goes deeper. In proving this, we calculate these functions using the core graph of H and its quotients. It turns out that both $\pi(H)$ and $\phi(H)$ can be computed explicitly via the subgraph of the DAG induced by all descendants of $\Gamma_X(H)$.

In the calculation of $\phi(H)$ we use the core graph $\Gamma_X(H)$ and its quotients to partition the event that 1 is a common fixed point of $\alpha_n(w)$ of each $w \in H$ (see Section 5).

Fortunately, the same core graph and quotients can also be used to find the primitivity rank $\pi(H)$, as shown in Section 4. Lemma 4.3 shows that all H-critical subgroups (see Definition 1.8) are always represented in the fringe (set of quotients) of H. Theorem 1.1 then shows directly how to calculate $\pi(H)$ using the fringe.

We show that under certain conditions, the two categorizations $\pi(\cdot)$ and $\phi(\cdot)$ indeed coincide.

Proposition 1.9. *Let $H \leq_{fg} \mathbf{F}_k$. Then for every $i \leq rk(H) + 1$,*

1. $\pi(H) = i \iff \phi(H) = i$
2. Moreover, if $\pi(H) = \phi(H) = i$ then $a_i(H)$ equals the number of H-critical subgroups of $\mathbf{F}_k$.

The second part of this proposition is in fact a generalization of a result of Nica. For a single element $w \in \mathbf{F}_k$ which is a proper power, namely $\pi(w) = \phi(w) = 1$, let $w = u^d$ with d maximal (so u is not a proper power). Let M denote the number of divisors of d. It is not hard to see that the number of w-critical subgroups of $\mathbf{F}_k$ equals $M - 1$: these subgroups are exactly $\langle u^m \rangle$ for every $1 \leq m < d$ such that $m|d$. This shows that the average number of fixed points in the permutation $\alpha_n(w)$ goes to M as $n \to \infty$. This corresponds to Corollary 1.3 in [Nic94] (for the case $L = 1$)[2].

[2]Nica's result was more general in a different manner: it involved the distribution of the number of L-cycles in the random permutation $\alpha_n(w)$, for any fixed L. He showed that as $n \to \infty$, the limit distribution depends only on d, where $w = u^d$ as above.

The connection between $\pi(\cdot)$ and $\phi(\cdot)$ goes beyond the cases stated in Proposition 1.9. To start off, if $\pi(H) = \infty$, then $H \overset{*}{\leq} \mathbf{F}_k$ and therefore H is measure preserving, and thus $\phi(H) = \infty$. In addition, Lemma 6.8 states that both $\pi(\cdot)$ and $\phi(\cdot)$ are additive with respect to concatenation of words on disjoint letter sets. Namely, if the words $w_1, w_2 \in \mathbf{F}_k$ have no letters in common then $\pi(w_1 w_2) = \pi(w_1) + \pi(w_2)$ and $\phi(w_1 w_2) = \phi(w_1) + \phi(w_2)$. Moreover, if the disjoint w_1 and w_2 satisfy both parts of Proposition 1.9 then so does their concatenation $w_1 w_2$.

In view of this discussion, the following conjecture suggests itself quite naturally:

Conjecture 1.10.

1. For every $H \leq_{fg} \mathbf{F}_k$
$$\pi(H) \quad = \quad \phi(H)$$

2. Moreover, $a_{\phi(H)}(H)$ equals the number of H-critical subgroups of $\mathbf{F}_k$.

Specifically, for a single word w, Proposition 1.9 states that for $i = 0, 1, 2$, $\pi(w) = i \Leftrightarrow \phi(w) = i$. As mentioned, the possible values of $\pi(H)$ are $\{0, 1, 2, \ldots, k\} \cup \{\infty\}$, and $\pi(H) = \infty$ iff $H \overset{*}{\leq} \mathbf{F}_k$. We also have $\pi(H) = \infty \Rightarrow \phi(H) = \infty$ (a free factor subgroup is measure preserving). Thus, when $rk(H) \geq k - 1$, the value of $\pi(H)$ uniquely determines $\phi(H)$ and the two values coincide. In other words, when $rk(H) \geq k - 1$
$$\pi(H) = \phi(H).$$

This shows, in turn, that when H is measure preserving, we have $\pi(H) = \phi(H) = \infty$, and so H is a free factor. This yields Theorem 1.5. The same argument shows that Conjecture 1.4 follows from part (1) of Conjecture 1.10 and suggests, in particular, a general strategy towards proving Conjecture 1.4.

As an aside, the second parts of Proposition 1.9 and Conjecture 1.10 say something interesting on the average number of fixed points in the random permutation $\alpha_n(w)$. We conjecture that for every w and for large enough n, this average is at least 1. In other words, among the family of distributions of S_n induced by free words, a random uniformly chosen permutation has the least average number of fixed points. This point is further elaborated in Section 8.

At this point we should clarify the relation of these results and some of what we did in [LP10]. There we introduced $\beta(\cdot)$ - yet another categorization of formal words. Just like $\phi(\cdot)$ and $\pi(\cdot)$ it maps every formal word to a non-negative integer or ∞. As it turns out, $\pi(\cdot)$ and $\beta(\cdot)$ coincide. This follows from Theorem 1.1 and from Section 4. The definition of $\pi(\cdot)$ is simpler and more elegant than the original definition of $\beta(\cdot)$. As shown in [LP10] for $i = 0, 1$, $\phi(w) = i \iff \beta(w) = i$. A partial proof was given there as well for the case $i = 2$. In Section 6 we complete the argument for

$i = 2$ and generalize it to prove Proposition 1.9.

The paper is arranged as follows. In section 2 we introduce the notions of core graphs, their morphisms and their quotients. In Section 3 we introduce our new perspective on core graphs, including the notion of immediate quotients and the mentioned DAG, and then prove Theorem 1.1. In Section 4 we analyze the primitivity rank of any $H \leq_{fg} \mathbf{F}_k$ and show how it can be computed from the quotients of $\Gamma_X(H)$ in the DAG of finite rank subgroups of $\mathbf{F}_k$. Section 5 is devoted to proving that $\phi(H)$ is well defined and can be indeed computed from the same descendants of $\Gamma_X(H)$. In Section 6 we establish the results connecting $\phi(\cdot)$ and $\pi(\cdot)$, culminating in the proof of Theorem 1.5. The concluding sections are devoted to two different consequences of the main results: the characterization of elements of $\mathbf{F}_k$ which are primitive in its profinite completion (Section 7) and the possible values of the average number of fixed points in the image of a word map on S_n (Section 8). The discussion in the three appendices is not necessary for the main results of this paper, but it does, in our view, complete the picture. In particular, we illustrate in Appendix A the algorithm to detect free factor subgroups.

2 Core Graphs and their Quotients

All groups that appear here are subgroups of $\mathbf{F}_k$, the free group with a given basis $X = \{x_1, \ldots, x_k\}$. Some of the relations we consider depend on the choice of the basis. We first describe core-graphs, which play a crucial role in this paper.

2.1 Core Graphs

Associated with every subgroup $H \leq \mathbf{F}_k$ is a directed, pointed, edge-labeled graph. This graph is called *the core-graph associated with* H and is denoted by $\Gamma_X(H)$. We recall the notion of $\overline{\Gamma}_X(H)$ the Schreier (right) coset graph of H with respect to the basis X. This is a directed, pointed and edge-labeled graph. Its vertex set is the set of all right cosets of H in $\mathbf{F}_k$, where the basepoint corresponds to the trivial coset H. For every coset Hw and every letter x_i there is a directed i-edge (short for x_i-edge) going from the vertex Hw to the vertex Hwx_i.

The core graph $\Gamma_X(H)$ is obtained from $\overline{\Gamma}_X(H)$ by omitting all the vertices and edges of $\overline{\Gamma}_X(H)$ which are never traced by a reduced (i.e., non-backtracking) path that starts and ends at the basepoint. Stated informally, we omit all (infinite) "hanging trees" from $\overline{\Gamma}_X(H)$. To illustrate, Figure 2.1 shows the graphs $\overline{\Gamma}_X(H)$ and $\Gamma_X(H)$ for $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle \leq \mathbf{F}_2$.

Note that the graph $\overline{\Gamma}_X(H)$ is $2k$-regular: Every vertex has exactly one outgoing j-edge and one incoming j-edge for every $1 \leq j \leq k$. Every vertex of $\Gamma_X(H)$ has *at most* one outgoing j-edge, and *at most* one incoming j-edge for every $1 \leq j \leq k$.

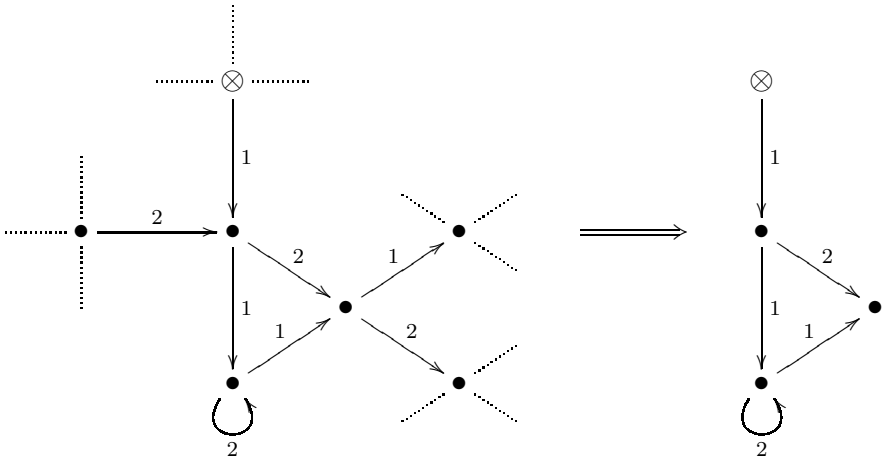

Figure 2.1: $\overline{\Gamma}_X(H)$ and $\Gamma_X(H)$ for $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle \leq \mathbf{F}_2$. The Schreier coset graph $\overline{\Gamma}_X(H)$ is the infinite graph on the left (the dotted lines represent infinite 4-regular trees). The basepoint "$\otimes$" corresponds to the trivial coset H, the vertex below it corresponds to the coset Hx_1, the one further down corresponds to $Hx_1^2 = Hx_1x_2x_1^{-1}$, etc. The core graph $\Gamma_X(H)$ is the finite graph on the right, which is obtained from $\overline{\Gamma}_X(H)$ by omitting all vertices and edges that are not traced by reduced closed paths around the basepoint.

It is an easy observation that

$$\pi_1(\overline{\Gamma}_X(H)) = \pi_1(\Gamma_X(H)) \overset{canonically}{\cong} H$$

where the canonical isomorphism is given by associating words in $\mathbf{F}_k$ to paths in the coset graph and in the core graph: We traverse the path by following the labels of outgoing edges. For instance, the path (from left to right)

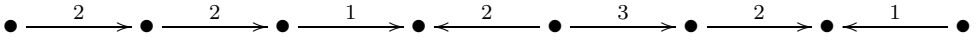

corresponds to the word $x_2^2 x_1 x_2^{-1} x_3 x_2 x_1^{-1}$. (See also [MVW07], where this fact appears in a slightly different language).

Core graphs were introduced by Stallings [Sta83]. Our definition is slightly different, in that we allow the basepoint to have degree one.

In fact, a "tail" in $\Gamma_X(H)$, i.e., a path to the basepoint can be eliminated by replacing H by an appropriate conjugate. However, we find it unnecessary and less elegant for our needs.

We now list some properties of core graph, most of which are proved in at least one of [Sta83, KM02, MVW07]. The remaining ones are easy observations.

Claim 2.1. *Let H be a subgroup of $\mathbf{F}_k$ with an associated core graph $\Gamma = \Gamma_X(H)$. The Euler Characteristic of a graph, denoted $\chi(\cdot)$ is the number of vertices minus the number of edges. Finally, $rk(H)$ denotes the rank of the group H.*

1. $rk(H) < \infty \Leftrightarrow \Gamma$ is finite
2. $rk(H) = 1 - \chi(\Gamma)$
3. Let Λ be a finite, pointed, directed graph with edges labeled by $\{x_1, \ldots, x_k\}$. Then Λ is a core graph (corresponding to some $J \leq \mathbf{F}_k$) if and only if Λ satisfies the following three properties:
   (a) Λ is connected
   (b) With the possible exception of the basepoint, every vertex has degree at least 2.
   (c) For every $1 \leq j \leq k$, no two j-edges share the same origin nor the same terminus.
4. There is a one-to-one correspondence between subgroups of $\mathbf{F}_k$ and core graphs.
5. There is a one-to-one correspondence between subgroups of $\mathbf{F}_k$ of finite rank and finite core graphs.

In Appendix C we present a well known algorithm, based on Stallings' foldings, to obtain the core graph of every $H \leq_{fg} \mathbf{F}_K$ given some finite generating set for H.

2.2 Morphisms of Core Graphs

In our framework, a morphism between two core-graphs Γ_1 and Γ_2 is a map that sends vertices to vertices and edges to edges, and preserves the structure of the graphs. Namely, it preserves the incidence relations, sends the basepoint to the basepoint, and preserves the directions and labels of the edges.

As in Claim 2.1, the proofs of the following properties are either easy variations on proofs in [Sta83, KM02, MVW07] or just easy observations:

Claim 2.2. *Let $H_1, H_2 \leq \mathbf{F}_k$ be subgroups, and Γ_1, Γ_2 be the corresponding core graphs. Then*

1. A morphism $\eta : \Gamma_1 \to \Gamma_2$ exists $\Leftrightarrow H_1 \leq H_2$,
   and in this case, $\eta_* : \pi_1(\Gamma_1) \to \pi_1(\Gamma_2)$ is injective.
2. If a morphism exists, it is unique.
3. Every morphism in an immersion (locally injective at the vertices).

2.3 Quotients of Core Graphs

With core-graph morphisms at hand, we can define the following rather natural relation between core-graphs.

Definition 2.3. Let Γ_1, Γ_2 be core graphs and $H_1, H_2 \leq \mathbf{F}_k$ the corresponding subgroups. We say that Γ_1 **covers** Γ_2 or that Γ_2 is a **quotient** of Γ_1 if there is a surjective morphism $\eta : \Gamma_1 \twoheadrightarrow \Gamma_2$. We also say in this case that H_1 **covers** H_2, and denote $\Gamma_1 \twoheadrightarrow \Gamma_2$ or $H_1 \overset{X}{\twoheadrightarrow} H_2$.

By "surjective" we mean surjective on both the vertices and the edges. Note that we use the term "covers" even though this is *not* a covering map in general (the morphism from Γ_1 to Γ_2 is always locally injective at the vertices, but not necessarily locally bijective).

For instance, $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle \leq \mathbf{F}_k$ covers the group $J = \langle x_2, x_1^2, x_1 x_2 x_1 \rangle$, the corresponding core graphs of which are the leftmost and rightmost graphs in Figure 2.2. As another example, every core graph Γ that contains edges of all labels covers the wedge graph Δ_k.

We already know (Claim 2.2) that if $H_1 \overset{X}{\twoheadrightarrow} H_2$ then, in particular, $H_1 \leq H_2$. However, the converse is incorrect. For example, the group $K = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2}, x_2 \rangle$ contains $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle$ (we simply added x_2 as a third generator), yet K is not a quotient of H: the morphism $\eta : \Gamma_X(H) \to \Gamma_X(K)$ does not contain the 2-loop at the basepoint of $\Gamma_X(K)$ in its image.

Note also that the relation $H_1 \overset{X}{\twoheadrightarrow} H_2$ *depends on the given generating set X of* $\mathbf{F}_k$. For example, if $H = \langle x_1 x_2 \rangle$ then $H \overset{X}{\twoheadrightarrow} \langle x_1, x_2 \rangle = F_2$. However, $x_1 x_2$ is primitive and could be taken as part of the original basis of $\mathbf{F}_2$. In that case, the core graph of H would consist of a single vertex and single loop and would have no quotients except for itself.

It is also interesting to note that every quotient of the core-graph Γ corresponds to some partition of $V(\Gamma)$ (the partition determined by the fibers of the morphism). We can simply draw a new graph with a vertex for each block in the partition, and a j-edge from block b_1 to block b_2 whenever there is some j-edge (v_1, v_2) in Γ_1 with $v_1 \in b_1, v_2 \in b_2$. However, not every partition of $V(\Gamma)$ corresponds to a quotient core-graph: In the resulting graph two distinct j-edges may have the same origin or the same terminus. Note that even if a partition P of $V(\Gamma)$ yields a quotient which is not a core-graph, this can be remedied. We can activate the folding process exemplified in Appendix C and obtain a core graph. The resulting partition P' of $V(\Gamma)$ is the finest partition which yields a quotient core-graph and which is still coarser than P. We illustrate this in Figure 2.2.

Lemma 2.4. *Every finite core-graph has a finite number of quotients.*
Equivalently, every $H \leq_{fg} \mathbf{F}_k$ covers a finite number of other subgroups.

Proof. The number of quotients of Γ is bounded from above by the number of partitions of $V(\Gamma)$. □

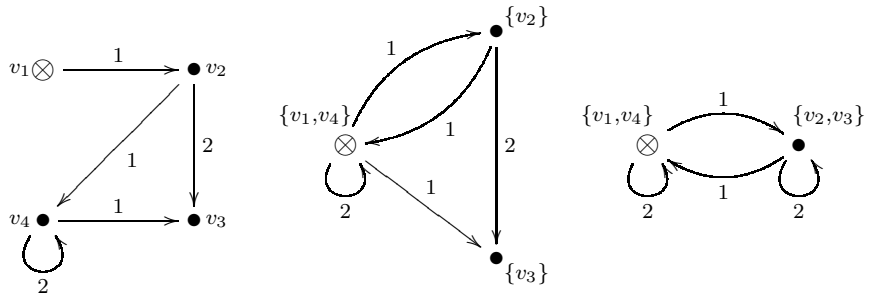

Figure 2.2: The left graph is the core graph $\Gamma_X(H)$ of $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle \leq \mathbf{F}_2$. Its vertices are denoted $v_1, \ldots, v_4$. The graph in the middle is the quotient corresponding to the partition $P = \{\{v_1, v_4\}, \{v_2\}, \{v_3\}\}$. This is not a core graph as there are two different 1-edges originating at $\{v_1, v_4\}$. In order to obtain a core quotient-graph, we use the folding process illustrated in Appendix C. The resulting core graph is on the right, corresponding to the partition $P' = \{\{v_1, v_4\}, \{v_2, v_3\}\}$.

Following the notations in [MVW07], we call the set of X-quotients of H the *X-fringe* of H and denote $\mathcal{O}_X(H)$. Namely,

$$\mathcal{O}_X(H) := \{\Gamma_X(J) \mid H \overset{X}{\twoheadrightarrow} J\} \tag{2.1}$$

Lemma 2.4 states in this terminology that for every $H \leq_{fg} \mathbf{F}_k$ (and every basis X), $|\mathcal{O}_X(H)| < \infty$.

Before describing our new perspective on core graphs, we remind some useful facts about free factors in free groups:

Claim 2.5. *Let $H, J, K \leq \mathbf{F}_k$. Then,*

1. Free factorness is transitive: If $H \overset{*}{\leq} J \overset{*}{\leq} K$ then $H \overset{*}{\leq} K$.
2. If $\eta : \Gamma_X(H) \hookrightarrow \Gamma_X(J)$ is an embedding then $H \overset{*}{\leq} J$.
3. If $H \overset{*}{\leq} J$ then H is a free factor in any subgroup $H \leq M \leq J$ in between.

Proof. The first and second claims are immediate. We give a "graph-theoretic" proof for the third one. Assume that $H \overset{*}{\leq} J$, and let Y be a basis of J extending some basis of H. In particular, $\Gamma_Y(H)$ and $\Gamma_Y(J)$ are both bouquets, consisting of a single vertex and $rk(H)$ (resp. $rk(J)$) loops. Now, for every $H \leq M \leq J$, consider the morphism $\eta : \Gamma_Y(H) \to \Gamma_Y(M)$. It is easy to see that a core-graph-morphism of a bouquet must be an embedding. Thus, by the second claim, $H \overset{*}{\leq} M$. □

3 Immediate Quotients and the DAG of Core Graphs

The quotient relation yields a partial order on the set of core graphs. But we are interested in a special case which we call *immediate quotients*. This relation allows us to build the aforementioned DAG (directed acyclic graph) of all (core graphs corresponding to) finite rank subgroups of $\mathbf{F}_k$.

Let Γ be a core graph, and let P be a partition of $V(\Gamma)$. Let Δ be the quotient core graph we obtain from P by the folding process described in Figures C.1 and 2.2. We say that Δ is *generated* from Γ by P. We are interested in the case where P identifies only a single pair of vertices:

Definition 3.1. Let Γ be a core graph and let P be a partition of $V(\Gamma)$ in which all parts consist of a single vertex with a single exceptional part that contains two vertices. Let Δ be the core graph generated by P. We then say that Δ is an **immediate quotient** of Γ.

Alternatively we say that Δ is *generated by merging a single pair* of vertices of Γ. For instance, the rightmost core graph in Figure 2.2 is an immediate quotient of the leftmost core graph.

The relation of immediate quotients has an interesting interpretation for the associated free groups. Let $H, J \leq \mathbf{F}_k$ be free groups and $\Gamma = \Gamma_X(H), \Delta = \Gamma_X(J)$ their core graph, and assume Δ is an immediate quotient of Γ obtained by identifying the vertices $u, v \in V(\Gamma)$. Now let $p_u, p_v \in \mathbf{F}_k$ be words corresponding to some paths in Γ from the basepoint to u and v respectively. It is not hard to see that identifying u and v has the same effect as adding the word $w = p_u p_v^{-1}$ to H and considering the generated group. Namely, $J = \langle H, w \rangle$.

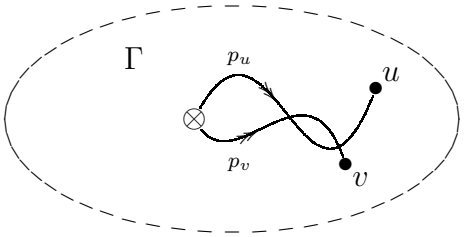

Based on the relation of immediate quotients we consider the DAG $\mathcal{D}_k$. The set of vertices of this graph consists of all finite core graphs with edges labeled by $1, \ldots, k$, and its directed edges connect every core graph to its immediate quotients. Every fixed ordered basis of $\mathbf{F}_k$ $X = \{x_1, \ldots, x_k\}$, determines a one-to-one correspondence between the vertices of this graph and all finite rank subgroups of $\mathbf{F}_k$.

As before, we fix an ordered basis X. For any $H \leq_{fg} \mathbf{F}_k$, the subgraph of $\mathcal{D}_k$ of the descendants of $\Gamma_X(H)$ consists of all quotients of $\Gamma_X(H)$, that is of all elements of the X-fringe $\mathcal{O}_X(H)$. By Lemma 2.4, this subgraph is finite. In Figure 3.1 we draw the subgraph of $\mathcal{D}_k$ consisting of all quotients of $\Gamma_X(H)$ when $H = \langle x_1 x_2 x_1^{-1} x_2^{-1} \rangle$.

The edges of this subgraph (i.e. immediate quotients) are denoted by the broken arrows in the figure.

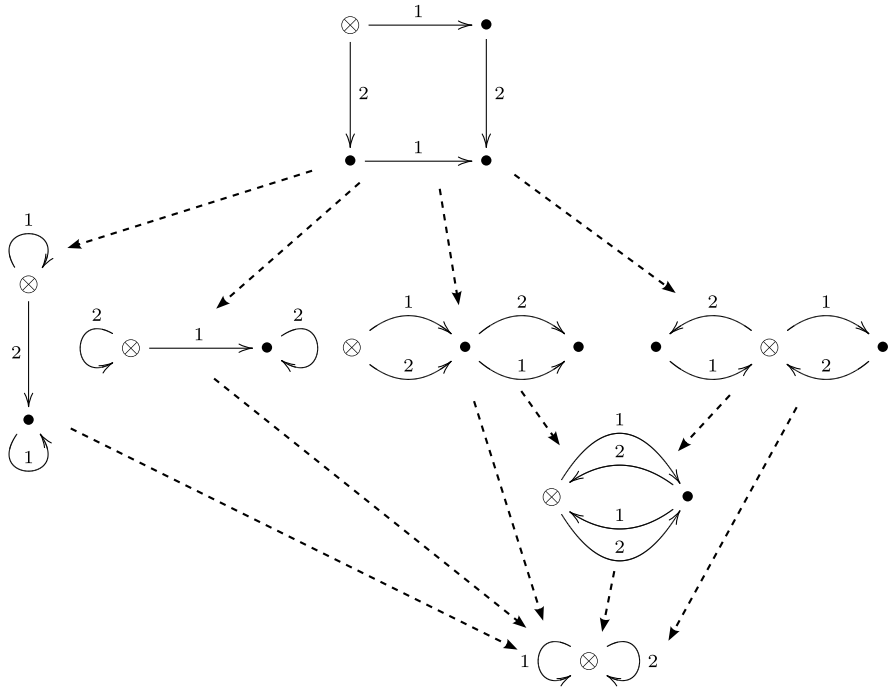

Figure 3.1: The subgraph of $\mathcal{D}_k$ induced by $\mathcal{O}_X(H)$, that is, all quotients of the core graph $\Gamma = \Gamma_X(H)$, for $H = \langle x_1 x_2 x_1^{-1} x_2^{-1} \rangle$. The dashed arrows denote immediate quotients, i.e. quotients generated by merging a single pair of vertices. Γ has exactly seven quotients: itself, four immediate quotients, and two quotients at distance 2.

This yields the aforementioned distance function between a finite core graph and a quotient of it:

Definition 3.2. Let $H_1, H_2 \leq_{fg} \mathbf{F}_k$ be finite rank subgroups such that $H_1 \overset{X}{\twoheadrightarrow} H_2$, and let $\Gamma_1 = \Gamma_X(H_1), \Gamma_2 = \Gamma_X(H_2)$ be the corresponding core graphs. We define the distance between H_1 and H_2, denoted $\rho_X(H_1, H_2)$ or $\rho(\Gamma_1, \Gamma_2)$, to be the shortest length of a directed path from Γ_1 to Γ_2 in $\mathcal{D}_k$.

In other words, $\rho_X(H_1, H_2)$ is the length of the shortest series of immediate quotients that yields Γ_2 from Γ_1. Equivalently, it is the minimal number of pairs of vertices that need to be identified in Γ_1 in order to obtain Γ_2 (via the folding process). For example, if Γ_2 is an immediate quotient of Γ_1 then $\rho_X(H_1, H_2) = \rho(\Gamma_1, \Gamma_2) = 1$. For $H = \langle x_1 x_2 x_1^{-1} x_2^{-1} \rangle$, $\Gamma_X(H)$ has four quotients at distance 1 and two at distance 2 (see Figure 3.1).

As aforementioned, by merging a single pair of vertices of $\Gamma_X(H)$ (and then folding) we obtain the core graph of a subgroup J obtained from H by adding some single generator. Thus, by taking an immediate quotient, the rank of the

associated subgroup increases at most by 1 (in fact, it may also stay unchanged or even decrease). This implies that whenever $H \overset{X}{\twoheadrightarrow} J$:

$$rk(J) - rk(H) \quad \leq \quad \rho(H, J) \tag{3.1}$$

It is not hard to bound the distance from above as well:

Lemma 3.3. *Let $H, J \leq_{fg} \mathbf{F}_k$ such that $H \overset{X}{\twoheadrightarrow} J$. Then*

$$rk(J) - rk(H) \quad \leq \quad \rho_X(H, J) \quad \leq \quad rk(J)$$

We postpone the proof of the upper bound to Appendix B. (In fact, this upper bound in not needed for the main results of this paper. We give it anyway in order to have the full picture in mind.) Theorem 1.1 then states that in the same setting, the lower bound is attained iff H is a free factor of J. In fact one can visualize these results in the following way. Consider a two dimensional table which contains all the elements of the fringe $\mathcal{O}_X(H)$ (each quotient of $\Gamma_X(H)$ contained in some, not necessarily distinct, cell). The rows correspond to the rank and are indexed $0, 1, 2, 3, \ldots$. The columns correspond to the distance from H and are also indexed $0, 1, 2, 3, \ldots$. We then put every quotient of H in the suitable cell in the table. Let $r = rk(H)$ denote the rank of H. Lemma 3.3 then says that the (finitely many) elements of $\mathcal{O}_X(H)$ are spread across $r+1$ diagonals in the table: the main one and the r diagonals below it. Theorem 1.1 implies that within $\mathcal{O}_X(H)$, H is a free factor of exactly those J-s found in the lowest of these $r+1$ diagonals. (In fact, Lemma 4.3 shows that $\pi(H)$ can also be read from this table: it equals the rank of the upmost occupied cell in this table outside the free-factor-diagonal.)

3.1 Proof of Theorem 1.1

The main result of this section states that if $H \leq_{fg} J \leq_{fg} \mathbf{F}_k$ and $H \overset{X}{\twoheadrightarrow} J$, then

$$\rho_X(H, J) = rk(J) - rk(H) \quad \Longleftrightarrow \quad H \overset{*}{\leq} J. \tag{3.2}$$

In fact, one of the implications is trivial. As mentioned above, merging two vertices in $\Gamma_X(H)$ is equivalent to adding some generator to H. If we manage to obtain $\Gamma_X(J)$ from $\Gamma_X(H)$ by $\mathrm{rk}(J) - \mathrm{rk}(H)$ merging steps, this means we can obtain J from H by adding $rk(J) - rk(H)$ extra generators to H, hence clearly $H \overset{*}{\leq} J$ (recall that by hopfianity of the free group, every generating set of size $\mathrm{rk}\,(J)$ is a basis of J, see e.g. [Bog08, Chapter 2.29]). Thus,

$$\rho_X(H, J) = rk(J) - rk(H) \quad \Longrightarrow \quad H \overset{*}{\leq} J \tag{3.3}$$

The other implication is not trivial. Assume that $H \overset{*}{\leq} J$. Our goal is to obtain $rk(J) - rk(H)$ complementary generators of J from H, so that each of them can be realized by merging a pair of vertices in $\Gamma_X(H)$.

To this goal we introduce the notion of a "*handle number*" associated with a subgroup M and a word $w \in \mathbf{F}_k$. (It also depends on the fixed basis X of $\mathbf{F}_k$). This number is defined as follows. Let $\Gamma = \Gamma_X(M)$. Denote by p_w the longest prefix of w that corresponds to some path from the basepoint of Γ (we trace the letters of w along Γ until we get stuck). Likewise, denote by s_w the longest suffix of w that ends at the basepoint (here we trace w^{-1} from the basepoint until we get stuck). If $|p_w| + |s_w| < |w|$, then $w = p_w m_w s_w$ as a reduced expression for some $1 \neq m_w \in \mathbf{F}_k$. The handle number of (M, w) is then

$$h_X(M, w) = h(\Gamma, w) = \begin{cases} |m_w| & |p_w| + |s_w| < |w| \\ 0 & \text{otherwise} \end{cases}.$$

Claim 3.4. *Assume that $w \notin M$ and let $N = \langle M, w \rangle$. Then*

1. *$h_X(M, w) > 0$ if and only if $\Gamma_X(M)$ is a (proper) subgraph of $\Gamma_X(N)$, and*
2. *$h_X(M, w) = 0$ if and only if $\Gamma_X(N)$ is an immediate quotient of $\Gamma_X(M)$.*

Proof. Assume first that $h_X(M, w) > 0$. In the notations of the previous paragraph, let $v(p_w), v(s_w)$ be the end point of the path corresponding to p_w and the starting point of the path corresponding to s_w. We can then add a "handle" to Γ in the form of a path corresponding to m_w which starts at $v(p_w)$ and ends at $v(s_w)$. (If $v(p_w) = v(s_w)$ this handle looks like a balloon, possibly with a string.)

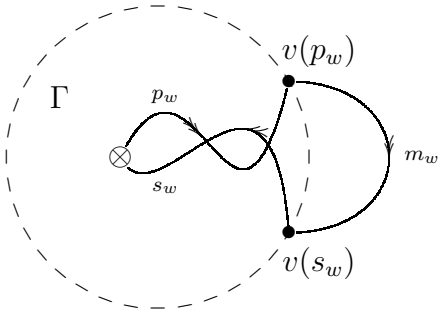

The resulting graph is a core-graph (the edge conditions at $v(p_w)$ and $v(s_w)$ hold, by the maximality of p_w and s_w), and it corresponds to N. So we actually obtained $\Gamma_X(N)$. It follows that $\Gamma_X(M)$ is a proper subgraph of $\Gamma_X(N)$. On the other hand, if $h_X(M, w) = 0$, i.e. if $|p_w| + |s_w| \geq |w|$, we can find a pair of vertices in Γ whose merging adds w to H as a complementary generator for J. (We may take $v(p_w)$ together with the vertex on the path of s_w at distance $|p_w| + |s_w| - |w|$ from $v(s_w)$.) □

The last claim shows, in particular, that if N is obtained from M by adding a single complementary generator, then either $\Gamma_X(N)$ is an immediate quotient or it contains $\Gamma_X(M)$ as a proper subgraph. This already proves Theorem 1.1 for the case

$\mathrm{rk}(J) - \mathrm{rk}(H) = 1$: if $H \overset{X}{\twoheadrightarrow} J$, we are clearly in the second case of Claim 3.4, i.e. J is an immediate quotient of H.

We proceed by defining handle numbers for a subgroup $M \leq_{fg} \mathbf{F}_k$ and an ordered set of words $w_1, \ldots, w_t \in \mathbf{F}_k$. Let $N = \langle M, w_1, \ldots, w_t \rangle$ and $\Delta = \Gamma_X(N)$. Let in addition $N_i = \langle M, w_1, \ldots, w_i \rangle$ and $\Gamma_i = \Gamma_X(N_i)$. We obtain a series of subgroups

$$M = N_0 \leq N_1 \leq \ldots \leq N_t = N,$$

and a series of graphs $\Gamma = \Gamma_0, \Gamma_1, \ldots, \Gamma_t = \Delta$. We denote by $h_X(M, w_1, \ldots, w_t)$ the t-tuple of the following handle numbers:

$$h_X(M, w_1, \ldots, w_t) \overset{\text{def}}{=} (h(\Gamma_0, w_1), h(\Gamma_1, w_2), \ldots, h_X(\Gamma_{t-1}, w_t)).$$

Let us focus now on the case where t is the cardinality of the smallest set $S \subseteq \mathbf{F}_k$ such that $N = \langle M, S \rangle$. The following lemma characterizes t-tuples of words for which the t-tuple of handle-numbers is lexicographically minimal. It is the crux of the proof of Theorem 1.1.

Lemma 3.5. *In the above notations, let $(w_1, \ldots, w_t)$ be an ordered set of complementary generators such that the tuple of handle numbers $h_X(M, w_1, \ldots, w_t)$ is lexicographically minimal. Then the zeros in $h_X(M, w_1, \ldots, w_t)$ form a prefix of the tuple.*

Namely, there is no zero handle-number that follows a positive handle-number.

Proof. It is enough to prove the claim for pairs of words (i.e. for $t = 2$), the general case following immediately. Assume then that $N = \langle M, w_1, w_2 \rangle$, that 2 is the minimal number of complementary generators of N given M, and that $h_X(M, w_1, w_2)$ is lexicographically minimal. In the above notation, assume to the contrary that $h(\Gamma_0, w_1) > 0$ and $h(\Gamma_1, w_2) = 0$. Let $m_1 = m_{w_1}$ denote the handle of w_1 in Γ_0. Thus Γ_1 was obtained from Γ_0 by adding a handle (or a balloon) representing m_1. The word w_2 can be expressed as $w_2 = ps$ so that there is a path corresponding to p in Γ_1, emanating from the basepoint and ending at some vertex u, and there is a path s to the basepoint from a vertex v. (Clearly, $u \neq v$ for otherwise $w_2 \in N_1$ contradicting the minimality of $t = 2$.) Thus Γ_2 is attained from Γ_1 by identifying the vertices u and v. By possibly multiplying w_2 from the left by a suitable element of N_1, we can assume that p does not traverse the handle m_1 "more than necessary". Namely, if u does not lie on m_1, then p avoids m_1, and if it does lie on m_1, then only the final segment of p traverses m_1 till u. The same holds for s and v (with right multiplication).

The argument splits into three possible cases.

- If both u, v belong already to Γ_0, then $h(\Gamma_0, w_2) = 0$. In this case we can switch between w_2 and w_1 to lexicographically reduce the sequence of handle numbers, contrary to our assumption.

- Consider next the case where, say, $v \in V(\Gamma_0)$ but $u \in V(\Gamma_1) \setminus V(\Gamma_0)$, i.e, u resides on the handle m_1. Then, the handle needed in order to add w_2 to Γ_0 is strictly shorter than $h(\Gamma_0, w_1) = |m_1|$. Again, by switching w_2 with w_1 the sequence of handle numbers goes down lexicographically - a contradiction.

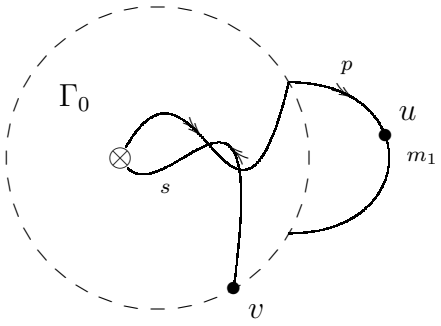

- In the final case that should be considered both u and v are on the handle m_1. I.e. $u, v \in V(\Gamma_1) \setminus V(\Gamma_0)$. Assume w.l.o.g. that when tracing the path of m_1, u precedes v. As before we can premultiply and postmultiply w_2 by suitable elements of N_1 to guarantee the following: The path p, from the basepoint of Γ_1 to u, goes through Γ_0 and then traverses a prefix of m_1 until reaching u, and the path s from v to the basepoint traces a suffix of m_1 and then goes only through Γ_0. Again $h(\Gamma_0, w_2) < h(\Gamma_0, w_1)$, so that switching w_2 with w_1 lexicographically reduces the sequence of handle numbers. (A similar argument works in the case m_1 constitutes a balloon.)

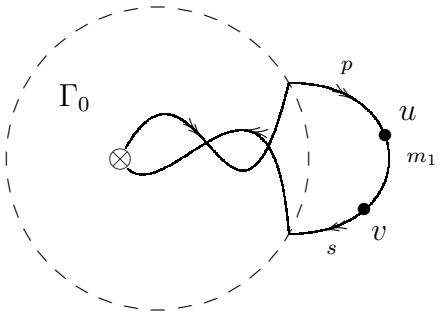

□

Theorem 1.1 follows easily from this lemma. Indeed, assume that $H \overset{*}{\leq} J$ and that $H \overset{X}{\twoheadrightarrow} J$. Let $t = \mathrm{rk}\,(J) - \mathrm{rk}\,(H)$ denote the difference in ranks, so that t is the smallest number of complementary generators needed to obtain J given H. Let $(w_1, \ldots, w_t)$ be an ordered set of complementary generators so that $h_X\,(H, w_1, \ldots, w_t)$ is lexicographically minimal. Similarly to the notations above, let $J_i = \langle H, w_1, \ldots, w_i \rangle$ and $\Gamma_i = \Gamma_X\,(J_i)$.

By the lemma, there is some $0 \leq q \leq t$ so that $h(\Gamma_0, w_1) = \ldots = h(\Gamma_{q-1}, w_q) = 0$ whereas $h(\Gamma_q, w_{q+1}), \ldots, h(\Gamma_{t-1}, w_t)$ are all positive. By Claim 3.4 it follows that Γ_i is an immediate quotient of Γ_{i-1} for $1 \leq i \leq q$, and therefore $J_q \in \mathcal{O}_X(H)$ and $\rho_X(H, J_q) = q$. (This in fact shows that $\rho_X(H, J_q) \leq rk(J_q) - rk(H)$, and the equality follows from Lemma 3.3).

Using Claim 3.4 again, we see that Γ_i is a proper subgraph of Γ_{i+1} for $q \leq i \leq t-1$. So that Γ_q is a subgraph of $\Gamma_t = \Gamma_X(J)$. But then the image of the graph morphism $\eta : \Gamma_X(H) \to \Gamma_X(J)$ is clearly the subgraph Γ_q. If $q < t$ this is a proper subgraph, which contradicts the assumption $H \overset{X}{\twoheadrightarrow} J$. Hence $q = t$ and $\rho_X(H, J) = t$, as required. Together with (3.3) this completes the proof of Theorem 1.1. □

In fact, the same argument yields a more general result:

Corollary 3.6. *Let $H \leq J \leq \mathbf{F}_k$ be f.g. groups, and let t be the minimal number of complementary generators needed to obtain J from H. Then t is computable as follows. Let $\eta : \Gamma_X(H) \to \Gamma_X(J)$ be the unique core-graph morphism, and let M be the intermediate subgroup corresponding to the image $\eta(\Gamma_X(H))$. Then,*

$$t = \rho_X(H, M) + \mathrm{rk}(J) - \mathrm{rk}(M).$$

Proof. In the notation of the last part of the proof of Theorem 1.1, we see that $M = J_q \overset{*}{\leq} J$, and so $\rho_X(H, M) + \mathrm{rk}(\mathrm{J}) - \mathrm{rk}(M) = \rho_X(H, J_q) + (\mathrm{t} - \mathrm{q}) = t$. □

Remark 3.7. Note that in the crucial arguments of the proof of Theorem 1.1 we did not use the fact that the groups where of finite rank. Indeed, this result can be carefully generalized to subgroups of $\mathbf{F}_k$ of infinite rank.

Remark 3.8. Another way to interpret Theorem 1.1 is by saying that if $H \overset{*}{\leq} J$ and $H \overset{X}{\twoheadrightarrow} J$ with $t = \mathrm{rk}(\mathrm{J}) - \mathrm{rk}(\mathrm{H})$, then there exists some set $\{w'_1, \ldots, w'_t\}$ of complementary generators such that each w_i can be realized by merging a pair of vertices in $\Gamma_X(H)$. To see this, let $w_1, \ldots, w_t$ be as in the proof above, so w_i can be realized by merging a pair of vertices u and v in Γ_{i-1}. Let $\eta_{i-1} : \Gamma_X(H) \to \Gamma_{i-1}$ be the surjective morphism, and pick any vertices in the fibers $u' \in \eta^{-1}(u)$, $v' \in \eta^{-1}(v)$. Let w'_i be some word corresponding to the merging of u' and v' in $\Gamma_X(H)$. It is not hard to see that for each i, $\langle H, w_1, \ldots, w_i \rangle = \langle H, w'_1, \ldots, w'_i \rangle$.

4 More on the Primitivity Rank

Recall Definitions 1.7 and 1.8 where we defined $\pi(w)$, the primitivity rank of a word $w \in \mathbf{F}_k$, and $\pi(H)$, the primitivity rank of $H \leq_{fg} \mathbf{F}_k$. In this subsection we prove some characteristics of this categorization of formal words, and show it actually depends only on the quotients of the core graph $\Gamma_X(H)$ (or $\Gamma_X(\langle w \rangle)$). The claims are stated for subgroups, and can be easily interpreted for elements with the usual correspondence between the element $w \neq 1$ and the subgroup it generates $\langle w \rangle$. We begin by characterizing the possible values of $\pi(H)$.

Lemma 4.1. *Let $H \leq_{fg} \mathbf{F}_k$. Then*

$$H \overset{*}{\leq} \mathbf{F}_k \Leftrightarrow \pi(H) = \infty.$$

Proof. Recall that $\pi(H)$ is defined by the smallest rank of subgroups of $\mathbf{F}_k$ where H is contained but not as a free factor. If H is not a free factor of $\mathbf{F}_k$, then $\mathbf{F}_k$ itself is one such subgroup so that $\pi(H) \leq k < \infty$. If $H \overset{*}{\leq} \mathbf{F}_k$, Claim 2.5 shows it is a free factor in every other subgroup containing it. Thus, in this case $\pi(H) = \infty$. □

Corollary 4.2. *For every $H \leq_{fg} \mathbf{F}_k$*

$$\pi(H) \in \{0, 1, \ldots, k\} \cup \{\infty\}$$

In the definition of the primitivity rank of a subgroup H, we consider all subgroups of $\mathbf{F}_k$ containing H but not as a free factor. It turns out it is enough to consider only subgroups of $\mathbf{F}_k$ that are covered by H, that is, groups whose associated core graphs are in the X-fringe $\mathcal{O}_X(H)$.

Lemma 4.3. *For every $H \leq_{fg} \mathbf{F}_k$*

$$\pi(H) = min \left\{ rk(J) \;\middle|\; \begin{array}{c} H \overset{X}{\twoheadrightarrow} J \;\; and \\ H \;\; is \;\; \textbf{\textit{not}} \;\; a \;\; free \;\; factor \;\; of \;\; J \end{array} \right\} \tag{4.1}$$

Moreover, all H-critical subgroups of $\mathbf{F}_k$ are covered by H.

Proof. Recall that H-critical subgroups of $\mathbf{F}_k$ are the subgroups of smallest rank in which H is not a free factor (so in particular their rank is exactly $\pi(H)$). It is enough to show that every H-critical subgroup has its associated core graph in the fringe $\mathcal{O}_X(H)$.

Consider an H-critical subgroup $J \leq \mathbf{F}_k$. This J contains H but not as a free factor. By Claim 2.2 there exists a morphism $\eta : \Gamma_X(H) \to \Gamma_X(J)$. If η is surjective then $H \overset{X}{\twoheadrightarrow} J$ and $\Gamma_X(J) \in \mathcal{O}_X(H)$. Otherwise, consider J', the group corresponding to the core graph $\eta(\Gamma_X(H))$. This graph, $\Gamma_X(J')$, is a strict subgraph of $\Gamma_X(J)$, and so $J' \overset{*}{\lneq} J$ (see Claim 2.5). In particular $H \overset{X}{\twoheadrightarrow} J'$ and $rk(J') < rk(J)$. It is impossible that $H \overset{*}{\leq} J'$, because by transitivity this would yield that $H \overset{*}{\leq} J$. Thus, J' is a subgroup in which H is a not free factor, and of smaller rank than J. This contradicts the fact that J is H-critical. □

We note that in the terminology of [KM02, MVW07], H-critical subgroups are merely a special kind of "algebraic extensions" of H. (An algebraic extension of H is a group J such that for every M with $H \leq M \lneq J$, M is not a free factor of J.) Specifically, H-critical subgroups are algebraic extensions of H of minimal rank, excluding H itself. Our proof actually shows the more general fact that all algebraic extensions of H can be found in the fringe (this fact appears in [KM02, MVW07]).

At this stage we can describe exactly how the primitivity rank of a subgroup $H \leq_{fg} \mathbf{F}_k$ can be computed. In fact, all algebraic extensions and critical subgroups of H can be immediately identified:

Corollary 4.4. *Consider the induced subgraph of $\mathcal{D}_k$ consisting of all core graphs in $\mathcal{O}_X(H)$. Then,*

- The algebraic extensions of H are precisely the core graphs which are not an immediate quotient of any other core graph of smaller rank.
- The H-critical subgroups are the algebraic extensions of smallest rank, excluding H itself, and $\pi(H)$ is their rank.

Proof. The second statement follows from the discussion above and from definition 1.8. The first statement holds trivially for H itself. If J is a proper algebraic extension of H, then by the proof of Lemma 4.3, $J \in \mathcal{O}_X(H)$. If $\Gamma_X(J)$ is an immediate quotient of some $\Gamma_X(M)$ of smaller rank, where $M \in \mathcal{O}_X(H)$, then $H \leq M \lneq J$ and by (the easier implication of) Theorem 1.1 we conclude $M \overset{*}{\leq} J$, a contradiction.

On the other hand, if $J \in \mathcal{O}_X(H)$ is not an algebraic extension of H, then there is some intermediate subgroup L such that $H \leq L \overset{*}{\lneq} J$. We can assume $L \in \mathcal{O}_X(H)$ for otherwise it can be replaced with L' corresponding to the image of the morphism $\eta : \Gamma_X(H) \to \Gamma_X(L)$ (whence $L' \in \mathcal{O}_X(H)$ and $H \leq L' \overset{*}{\leq} L \overset{*}{\lneq} J$). From (the harder implication of) Theorem 1.1 it follows that $\rho_X(L, J) = \mathrm{rk}(J) - \mathrm{rk}(L)$. The prior-to-last element in a shortest path in $\mathcal{D}_k$ from $\Gamma_X(L)$ to $\Gamma_X(J)$ is then a proper free factor of J at distance 1 that belongs to $\mathcal{O}_X(H)$. □

As an example, consider $H = \langle x_1 x_2 x_1^{-1} x_2^{-1} \rangle$. The full lattice of groups in $\mathcal{O}_X(H)$ is given in Figure 3.1. There is one group of rank 1 (H itself), 5 of rank 2 and one of rank 3. The only group in the lattice where H in not a free factor is the group $\langle x_1 x_2 \rangle = \mathbf{F}_2$, of rank 2, so $\pi(H) = 2$. (And the set of algebraic extensions of H is precisely $\{H, \mathbf{F}_2\}$.)

5 The Calculation of ϕ

The proof of Proposition 1.9, as well as the reasoning that underlies Conjecture 1.10, are based on the fact that both $\phi(H)$ and $\pi(H)$ can be calculated by analyzing $\mathcal{O}_X(H)$, the set of quotients of $\Gamma_X(H)$. In the previous section it was shown how $\pi(H)$ is determined from $\mathcal{O}_X(H)$. In this section we show how $\phi(H)$ can be calculated by a simple analysis of the very same set. The origins of the algorithm we present here can be traced to [Nic94] with further development in [LP10]. We present it here from a more general perspective.

Let the group G act on a set Y and let $y_0 \in Y$ be a fixed element. Consider a random homomorphism $\alpha_G \in Hom(F_k, G)$. The core graphs in $\mathcal{O}_X(H)$ can be used

to calculate the probability that $\alpha_G(H) \subset G_{y_0}$ (where G_{y_0} is the stabilizer of the element y_0). The quotients of the core graph $\Gamma_X(H)$ correspond to all the different "coincidence patterns" of the paths of y_0 through the action of the images of all $w \in H$, thereby describing disjoint events whose union is the event that $\alpha_G(H) \subset G_{y_0}$.

The idea is that in order to determine whether $\alpha_G(w)$ stabilizes y_0 for some $w \in \mathbf{F}_k$, we do not need to know all the values $\alpha_G(x_i)$ over $x_i \in X$ (the given basis of $\mathbf{F}_k$). Rather, we only need to know how $\alpha_G(x_i)$ acts on certain elements in Y, those in the path of y_0 through $\alpha_G(w)$. Namely, if $w = x_{j_1}^{\epsilon_1} \dots x_{j_{|w|}}^{\epsilon_{|w|}}$, $j_i \in \{1, \dots, k\}, \epsilon_i \in \{\pm 1\}$, we need to uncover the elements $y_1, \dots, y_{|w|}$ in the following diagram:

$$y_0 \xrightarrow{\alpha_G\left(x_{j_1}^{\epsilon_1}\right)} y_1 \xrightarrow{\alpha_G\left(x_{j_2}^{\epsilon_2}\right)} y_2 \xrightarrow{\alpha_G\left(x_{j_3}^{\epsilon_3}\right)} \cdots \xrightarrow{\alpha_G\left(x_{j_{|w|-1}}^{\epsilon_{|w|-1}}\right)} y_{|w|-1} \xrightarrow{\alpha_G\left(x_{j_{|w|}}^{\epsilon_{|w|}}\right)} y_{|w|}$$

That is, the image of $x_{j_1}^{\epsilon_1}$ acts on y_0, and we denote the resulting element by $y_1 \in Y$. The image of y_1 under the action of $\alpha_G\left(x_{j_2}^{\epsilon_2}\right)$ is denoted by y_2, etc. Then, y_0 is a fixed point of $\alpha_G(w)$ iff $y_{|w|} = y_0$.

There are normally many possible series of elements $y_1, \dots, y_{|w|-1} \in Y$ that can constitute the path of y_0 through $\alpha_G(w)$ such that y_0 is a fixed point. We divide these different series to a *finite* number of categories based on the *coincidence pattern* of this series. Namely, two realizations of this series, $y_1, \dots, y_{|w|-1}$, and $y'_1, \dots, y'_{|w|-1}$ are in the same category iff for every $i, j \in \{0, \dots, |w|-1\}$, $y_i = y_j \Leftrightarrow y'_i = y'_j$ (note that the elements of the series are also compared to y_0). In other words, every coincidence pattern corresponds to some partition of $\{0, 1, \dots, |w|-1\}$.

However, because the elements $\alpha_G(x_j) \in G$ act as permutations on Y, not every partition yields a realizable coincidence pattern: if, for example, $x_{j_2}^{\epsilon_2} = x_{j_7}^{-\epsilon_7}$, and $y_1 = y_7$, we must also have $y_2 = y_6$. This condition should sound familiar. Indeed, for each coincidence pattern we can draw a pointed, directed, edge-labeled graph describing it. The vertices of this graph correspond to blocks in the partition of $\{0, 1, \dots, |w|-1\}$, the basepoint corresponding to the block containing 0. Then, for each $i \in \{1, \dots, |w|\}$ there is a j_i-edge, directed according to ϵ_i, between the block of $i-1$ to the block of i. The constraints that coincidence patterns must satisfy then becomes the very same ones we had encountered in our discussion of core graphs. Namely, no two j-edges share the same origin or the same terminus.

Thus, the different realizable coincidence patterns of the series $y_0, y_1, \dots, y_{|w|-1}$ are exactly those described by core graphs that are quotients of $\Gamma_X(\langle w \rangle)$. For instance, there are exactly seven realizable coincidence patterns that correspond to the event in which y_0 is a fixed point of $\alpha_G(w)$ when $w = [x_1, x_2]$. The seven core graphs in Figure 3.1 correspond to these seven coincidence patterns.

Finally, the same phenomenon generalizes to any $H \leq_{fg} \mathbf{F}_k$. Instead of uncovering the path of y_0 through the image of a single word, we uncover the paths trough all words in H. The union of these paths in which y_0 is stabilized by all elements of H is depicted exactly by the core graph $\Gamma_X(H)$. The realizable coincidence patterns correspond then to the quotients of $\Gamma_X(H)$, namely to $\mathcal{O}_X(H)$. To summarize:

$$Prob\big[\alpha_G(H) \subset G_{y_0}\big] = \sum_{\Gamma \in \mathcal{O}_X(H)} Prob\big[\begin{smallmatrix}\Gamma \text{ describes the coincidence pattern} \\ \text{of } y_0 \text{ through the action of } \alpha_G(H)\end{smallmatrix}\big] \quad (5.1)$$

The advantage of the symmetric group S_n with its action on $\{1,\ldots,n\}$ is that the probabilities in the r.h.s. of (5.1) are very easy to formulate. Let $\alpha_n = \alpha_{S_n} \in Hom(\mathbf{F}_k, S_n)$ be a uniformly distributed random homomorphism, and let $\Gamma \in \mathcal{O}_X(H)$ be one of the quotients of $\Gamma_X(H)$. Denote by $P_\Gamma(n)$ the probability that $\alpha_n(H) \subset (S_n)_1$ and that the coincidence pattern of the paths of 1 through the elements $\alpha_G(H)$ are described by Γ. Then we can give an exact expression for $P_\Gamma(n)$ in terms of v_Γ, e_Γ and e^j_Γ, the number of vertices, edges and j-edges in Γ:

There are $(n-1)(n-2)\ldots(n-v_\Gamma+1)$ possible assignments of different elements from $\{2,3,\ldots,n\}$ to the vertices of Γ (excluding the basepoint which always corresponds to the element 1). Then, for a given assignment, there are exactly e^j_Γ constraints on the permutation $\alpha_n(x_j)$. So the probability that the permutation $\alpha_n(x_j)$ agrees with the given assignment is

$$\frac{(n-e^j_\Gamma)!}{n!} = \frac{1}{n(n-1)\ldots(n-e^j_\Gamma+1)}$$

(for $n \geq e^j_\Gamma$). Thus

$$P_\Gamma(n) = \frac{(n-1)(n-2)\ldots(n-v_\Gamma+1)}{\prod_{j=1}^k n(n-1)\ldots(n-e^j_\Gamma+1)}$$

Recall the definition of $\Phi_H(n)$ in (1.1). Since for every j and every $\Gamma \in \mathcal{O}_X(H)$ we have $e^j_\Gamma \leq e^j_{\Gamma_X(H)}$ we can summarize and say that for every $n \geq \max_j e^j_{\Gamma_X(H)}$ (in particular for every $n \geq v_{\Gamma_X(H)}$), we have:

$$\begin{aligned}
\Phi_H(n) &= Prob\left[\forall w \in H \quad \alpha_n(w)(1) = 1\right] - \frac{1}{n^{rk(H)}} \\
&= Prob\left[\alpha_n(H) \subset (S_n)_1\right] - \frac{1}{n^{rk(H)}} \\
&= -\frac{1}{n^{rk(H)}} + \sum_{\Gamma \in \mathcal{O}_X(H)} \frac{(n-1)(n-2)\ldots(n-v_\Gamma+1)}{\prod_{j=1}^k n(n-1)\ldots(n-e^j_\Gamma+1)} \\
&= -\frac{1}{n^{rk(H)}} + \sum_{\Gamma \in \mathcal{O}_X(H)} \frac{1}{n^{e_\Gamma - v_\Gamma + 1}} \frac{(1-\frac{1}{n})(1-\frac{2}{n})\ldots(1-\frac{v_\Gamma-1}{n})}{\prod_{j=1}^k (1-\frac{1}{n})\ldots(1-\frac{e^j_\Gamma-1}{n})} \qquad (5.2)
\end{aligned}$$

For instance, for $H = \langle [x_1, x_2] \rangle$ there are seven summands in the r.h.s. of (5.2), corresponding to the seven core graphs in Figure 3.1. If we go over these core graphs

from top to bottom and left to right, we obtain that for every $n \geq 2$:

$$\begin{aligned}\Phi_{\langle [x_1,x_2]\rangle}(n) = &-\frac{1}{n} + \frac{(n-1)(n-2)(n-3)}{n(n-1)\cdot n(n-1)} + \\ &+ \frac{n-1}{n(n-1)\cdot n} + \frac{n-1}{n\cdot n(n-1)} + \frac{(n-1)(n-2)}{n(n-1)\cdot n(n-1)} + \\ &+ \frac{(n-1)(n-2)}{n(n-1)\cdot n(n-1)} + \frac{n-1}{n(n-1)\cdot n(n-1)} + \frac{1}{n\cdot n} \\ = &-\frac{1}{n} + \frac{1}{n-1} = \frac{1}{n(n-1)}\end{aligned}$$

Recall the definition of the second categorization of sets of free words, $\phi(H)$, in (1.3). Indeed, we can rewrite (5.2) as a power series in $\frac{1}{n}$, and obtain that (for large enough n)

$$\Phi_H(n) = \sum_{i=0}^{\infty} \frac{a_i(H)}{n^i}$$

where the coefficients $a_i(H)$ depend only on H. We need not consider negative values of i because the leading term of every summand in (5.2) is $\frac{1}{n^{e_\Gamma - v_\Gamma + 1}}$, and $e_\Gamma - v_\Gamma + 1$ is non-negative for connected graphs. In fact, this number also equals the rank of the free subgroup corresponding to Γ.

The value of $\phi(H)$ equals the smallest i for which $a_i(H)$ does not vanish. For instance, for $H = \langle [x_1, x_2]\rangle$ we have

$$\Phi_{\langle [x_1,x_2]\rangle}(n) = \frac{1}{n(n-1)} = \sum_{i=2}^{\infty} \frac{1}{n^i}$$

so that $a_0(H) = a_1(H) = 0$ and $a_i(H) = 1$ for $i \geq 2$. Hence $\phi(H) = 2$.

In fact, we can write a power series for each $\Gamma \in \mathcal{O}_X(H)$ separately, and obtain:

$$\begin{aligned}P_\Gamma(n) &= \frac{1}{n^{e_\Gamma - v_\Gamma + 1}} \frac{(1-\frac{1}{n})(1-\frac{2}{n})\dots(1-\frac{v_\Gamma - 1}{n})}{\prod_{j=1}^{k}(1-\frac{1}{n})\dots(1-\frac{e_\Gamma^j - 1}{n})} \\ &= \frac{1}{n^{e_\Gamma - v_\Gamma + 1}}\big(1 - \frac{\binom{v_\Gamma}{2} - \sum_{j=1}^{k}\binom{e_\Gamma^j}{2}}{n} + O\big(\frac{1}{n^2}\big)\big) \qquad (5.3)\end{aligned}$$

This shows that if $\Gamma = \Gamma_X(J)$ ($J \leq_{fg} \mathbf{F}_k$), then $P_\Gamma(n)$ never affects $a_i(H)$-s with $i < rk(J)$. It is also easy to see that all the coefficients of the power series expressing $P_\Gamma(n)$ are integers. We summarize:

Claim 5.1. *For every $H \leq_{fg} \mathbf{F}_k$, all the coefficients $a_i(H)$ are integers. Moreover, $a_i(H)$ is completely determined by core graphs in $\mathcal{O}_X(H)$ corresponding to groups of rank $\leq i$.*

6 Relations between $\pi(\cdot)$ and $\phi(\cdot)$

We now have all the background needed for the proof of Proposition 1.9 and consequently of Theorem 1.5. We need to show that for every $H \leq_{fg} \mathbf{F}_k$ and every $i \leq rk(H)+1$, we have

$$\pi(H) = i \Longleftrightarrow \phi(H) = i.$$

The proof is divided into three steps. First we deal with the case $i < rk(H)$, then with $i = rk(H)$. The last case $i = rk(H)+1$ is by far the hardest.

Lemma 6.1. *Let $H \leq_{fg} \mathbf{F}_k$ and $i < rk(H)$. Then*

1. $\pi(H) = i \Leftrightarrow \phi(H) = i$
2. If $\pi(H) = \phi(H) = i$ then $a_i(H)$ equals the number of H-critical subgroups of $\mathbf{F}_k$.

Proof. Let m denote the smallest rank of a group $J \leq \mathbf{F}_k$ such that $H \overset{X}{\twoheadrightarrow} J$ (so $m \leq \mathrm{rk}\,(\mathrm{H})$). The first part of the result is derived from the observation that both $\pi(H) = i$ and $\phi(H) = i$ iff $m = i$. Let us note first that $\pi(H) = i \Leftrightarrow m = i$. This follows from Lemma 4.3 and the fact that H cannot be a free factor in a subgroup of smaller rank.

We next observe that $\phi(H) = i \Leftrightarrow m = i$: If $m < rk(H)$ then by (5.2) and (5.3), m is indeed the smallest index for which $a_m(H)$ does not vanish (this does not work for $m = rk(H)$ because of the term $\left(-\frac{1}{n^{rk(H)}}\right)$ in the definition of $\Phi_H(n)$). Conversely, if $m = rk(H)$ then obviously $\phi(H) \geq rk(H)$.

For the second part of the lemma, recall that H is not a free factor in any subgroup of smaller rank containing it. Thus, by (5.3) and Lemma 4.3, both $a_i(H)$ and the number of subgroups of rank i containing H equal the number of subgroups of rank i in $\mathcal{O}_X(H)$. □

The case $i = rk(H)$ is slightly different, but almost as easy.

Lemma 6.2. *Let $H \leq_{fg} \mathbf{F}_k$. Then,*

1. $\pi(H) = rk(H) \Leftrightarrow \phi(H) = rk(H)$
2. If $\pi(H) = \phi(H) = rk(H)$ then $a_{rk(H)}(H)$ equals the number of H-critical subgroups of $\mathbf{F}_k$.

Proof. From Lemma 6.1 we infer that $\pi(H) \geq rk(H) \Leftrightarrow \phi(H) \geq rk(H)$. So we assume that indeed $\pi(H), \phi(H) \geq rk(H)$, or, equivalently, that there are no subgroups covered by H of rank smaller than $rk(H)$.

We show that both sides of part (1) are equivalent to the existence of a quotient (corresponding to a subgroup) of rank $rk(H)$ in $\mathcal{O}_X(H)$ *other* than $\Gamma_X(H)$ itself.

Indeed, this is true for $\pi(H)$ because the only free product of H of rank $rk(H)$ is H itself.

As for $\phi(H)$, this is true because when $\phi(H) \geq rk(H)$ it is easily verified that the value of $a_{rk(H)}(H)$ equals the number of quotient in $\mathcal{O}_X(H)$ of rank $rk(H)$ minus 1 (this minus 1 comes from the term $\left(-\frac{1}{n^{rk(H)}}\right)$). We think of this term as offsetting the contribution of $\Gamma_X(H)$ to $a_{rk(H)}(H)$, so $a_{rk(H)}(H)$ equals the number of other quotients in $\mathcal{O}_X(H)$ of rank $rk(H)$.

The second part of the lemma is true because all H-critical subgroups are covered by H (Lemma 4.3). □

6.1 The Case $i = rk(H) + 1$

The most interesting (and the hardest) case of Theorem 1.5 is when $rk(H) = k - 1$. In the previous analysis this corresponds to $i = rk(H) + 1$.

Lemma 6.3. *Let $H \leq_{fg} \mathbf{F}_k$. Then,*

1. $\pi(H) = rk(H) + 1 \Leftrightarrow \phi(H) = rk(H) + 1$
2. If $\pi(H) = \phi(H) = rk(H) + 1$ then $a_{rk(H)+1}(H)$ equals the number of H-critical subgroups of $\mathbf{F}_k$.

Denote by $\hat{\Gamma} = \Gamma_X(H)$ the associated core graph. By Lemmas 6.1 and 6.2, we can assume that $\pi(H), \phi(H) \geq rk(H) + 1$. In particular, we can thus assume that H is not contained in any subgroup of rank smaller than $rk(H) + 1$ other than H itself.

The coefficient $a_{rk(H)+1}(H)$ in the expression of $\Phi_H(n)$ is the sum of two expressions:

- The contribution of $\hat{\Gamma}$ which equals $-\left(\binom{v_{\hat{\Gamma}}}{2} - \sum_{j=1}^{k}\binom{e^j_{\hat{\Gamma}}}{2}\right)$
- A contribution of 1 from each core graph of rank $rk(H) + 1$ in $\mathcal{O}_X(H)$

Thus, our goal is to show that the contribution of $\hat{\Gamma}$ is exactly offset by the contribution of the core graphs of rank $rk(H) + 1$ in $\mathcal{O}_X(H)$ in which H is a free factor. This would then yield immediately both parts of Lemma 6.3. But the number of subgroups of rank $rk(H) + 1$ (in $\mathcal{O}_X(H)$) in which H is a free factor equals exactly the number of immediate quotients of $\hat{\Gamma}$: Theorem 1.1 shows that only immediate quotients of $\hat{\Gamma}$ are subgroups of rank $rk(H) + 1$ in which H is a free factor. On the other hand, (3.1) and the assumption that H in not contained in any other subgroup of equal or smaller rank yield that every immediate quotient of $\hat{\Gamma}$ is of rank $rk(H)+1$ (and H is a free factor in it).

Thus, Lemma 6.3 follows from the following lemma.

Lemma 6.4. *Assume $\pi(H), \phi(H) > rk(H)$. Then $\hat{\Gamma} = \Gamma_X(H)$ has exactly*

$$\binom{v_{\hat{\Gamma}}}{2} - \sum_{j=1}^{k} \binom{e^j_{\hat{\Gamma}}}{2}$$

immediate quotients.

The intuition behind Lemma 6.4 is this: Every immediate quotient of $\hat{\Gamma}$ is generated by identifying some pair of vertices of $\hat{\Gamma}$, and there are exactly $\binom{v_{\hat{\Gamma}}}{2}$ such pairs. But for every pair of equally-labeled edges of $\hat{\Gamma}$, the pair of origins generates the same immediate quotient as the pair of termini. This intuition needs, however, some justification that we now provide.

To this end we use the graph Υ, a concept introduced in [LP10][3]. This graph represents the pairs of vertices of $\hat{\Gamma}$ and the equivalence relations between them induced by equally-labeled edges. There are $\binom{v_{\hat{\Gamma}}}{2}$ vertices in Υ, one for each unordered pair of vertices of $\hat{\Gamma}$. The number of directed edges in Υ is $\sum_{j=1}^{k} \binom{e^j_{\hat{\Gamma}}}{2}$, one for each pair of equally-labeled edges in $\hat{\Gamma}$. The edge corresponding to the pair $\{\epsilon_1, \epsilon_2\}$ of j-edges is a j-edge connecting the vertex $\{origin(\epsilon_1), origin(\epsilon_2)\}$ to $\{terminus(\epsilon_1), terminus(\epsilon_2)\}$. For example, when S consists of the commutator word, Υ has $\binom{4}{2} = 6$ vertices and $\binom{2}{2} + \binom{2}{2} = 2$ edges. We illustrate a slightly more interesting case in Figure 6.1.

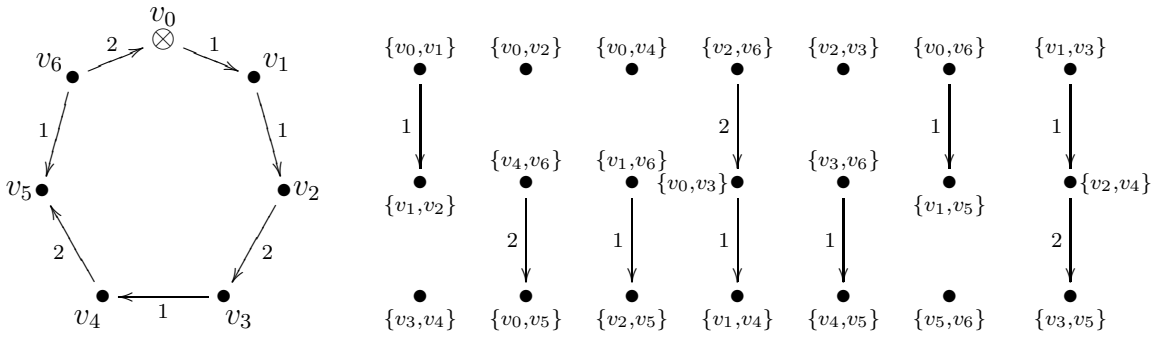

Figure 6.1: The graph Υ (on the right) corresponding to $\hat{\Gamma} = \Gamma_X(H)$ (on the left) for $H = \langle x_1^2 x_2 x_1 x_2 x_1^{-1} x_2 \rangle$. (The vertices of $\hat{\Gamma}$ are denoted here by $v_0, \ldots, v_6$.)

We denote the set of connected components of Υ by $Comp(\Upsilon)$. The proof of Lemma 6.4 will follow from two facts that we show next. Namely, Υ has exactly $\binom{v_{\hat{\Gamma}}}{2} - \sum_{j=1}^{k} \binom{e^j_{\hat{\Gamma}}}{2}$ connected components. Also, there is a one-to-one correspondence between $Comp(\Upsilon)$ and the set of immediate quotients of $\hat{\Gamma}$.

[3]This is a variation of the classical construction of *pull-back* of graphs (in this case the pull-back of the graph $\hat{\Gamma}$ with itself).

Claim 6.5. *If $\pi(H), \phi(H) > rk(H)$, then*

$$\left|Comp(\Upsilon)\right| = \binom{v_{\hat{\Gamma}}}{2} - \sum_{j=1}^{k} \binom{e^j_{\hat{\Gamma}}}{2}$$

Proof. Since Υ has $\binom{v_{\hat{\Gamma}}}{2}$ vertices and $\sum_{j=1}^{k} \binom{e^j_{\hat{\Gamma}}}{2}$ edges, it is enough to show that it is a forest, i.e., it contains no cycles.

Let $C \in Comp(\Upsilon)$ be some component of Υ. Clearly, every vertex in C (which corresponds to a pair of vertices in $\hat{\Gamma}$) generates the same immediate quotient. Denote this quotient by $\Delta(C)$, and the corresponding subgroup by J. In particular, $rk(J) = rk(H) + 1$ (recall that under the claim's assumptions, H is not contained in any other subgroup of smaller or equal rank).Assume to the contrary that C contains a cycle. Edges in Υ are directed and labeled, so there is an element $u \in \mathbf{F}_k$ which corresponds to this cycle started, say, at the vertex $\{x, y\}$.

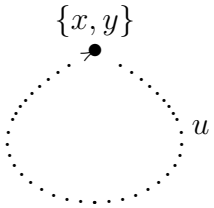

Where do we get as we walk in the core graph $\hat{\Gamma}$ starting at x (resp. y) and following the path corresponding to u? One possibility is that the walk from x returns back to x and likewise for y. Alternatively this u-walk can take us from x to y and from y to x. We consider only the former possibility. The latter case would be handled by considering the walk corresponding to u^2. Let $p_x, p_y \in \mathbf{F}_k$ be words corresponding to some paths from the basepoint of $\hat{\Gamma}$ to x, y respectively. In particular, $p_x u p_x^{-1}, p_y u p_y^{-1} \in H$.

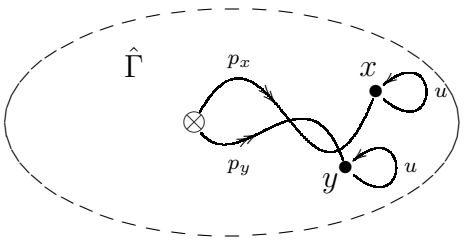

Merging x and y is equivalent to adding the generator $p_x p_y^{-1}$ to H, so that $J = \langle H, p_x p_y^{-1} \rangle$. Since $rk(J) = rk(H) + 1$, we have that $J = H * \langle p_x p_y^{-1} \rangle$. Consider the elements $h_1 = p_x u p_x^{-1} \in H$ and $h_2 = p_y u p_y^{-1} \in H$. The following equality holds:

$$h_1 = p_x u p_x^{-1} = (p_x p_y^{-1}) p_y u p_y^{-1} (p_x p_y^{-1})^{-1} = (p_x p_y^{-1}) h_2 (p_x p_y^{-1})^{-1}$$

This is a contradiction, since we obtained two different expressions for h_1 in the free product $J = H * \langle p_x p_y^{-1} \rangle$. □

Remark 6.6. Let $H \leq \mathbf{F}_k$ and $x, y \in V(\Gamma_X(H))$. In the proof of the last claim it was shown that if there is some $1 \neq u \in \mathbf{F}_k$ which is readable as a closed path at both x and y, then the subgroup we obtain by merging them is *not* a free extension of H. We stress that the converse is not true. For example, consider $H = \langle a, bab \rangle \leq \mathbf{F}(\{a, b\})$. Then $\Gamma_{\{a,b\}}(H)$ has three vertices, no pair of which share a common closed path (in other words, the corresponding graph Υ has no cycles). However, $\Gamma_{\{a,b\}}(H)$ has exactly two immediate quotients, none of which is a free extension.

Next we exhibit a one-to-one correspondence between $Comp(\Upsilon)$ and the immediate quotients of $\hat{\Gamma}$. It is very suggestive to try and restore C from $\Delta(C)$ by simply signaling out the pairs of vertices that are identified in $\Delta(C)$. But this is too naive. There may be pairs of vertices not in C that are identified in $\Delta(C)$. For instance, consider C, the rightmost component of Υ in Figure 6.1. In $\Delta(C)$ we merge v_1 and v_3 but also v_3 and v_5. Thus v_1 and v_5 are merged and likewise all pairs in the component of $\{v_1, v_5\}$.

However, simple group-theoretic arguments do yield this sought-after result:

Claim 6.7. *If $\pi(H), \phi(H) > rk(H)$, then there is a one-to-one correspondence between $Comp(\Upsilon)$ and the set of immediate quotients of $\hat{\Gamma} = \Gamma_X(H)$.*

Proof. Maintaining the above notation, we need to show that the mapping from $C \in Comp(\Upsilon)$ to $\Delta(C)$, the immediate quotient generated by any of the pairs in C, is one to one.

Denote by J the subgroup corresponding to the immediate quotient $\Delta(C)$. Let $\{x, y\}$ be some vertex in C, and $p_x, p_y \in \mathbf{F}_k$ words corresponding to some paths from the basepoint of $\hat{\Gamma}$ to x, y, respectively. Let also $q = p_x p_y^{-1} \in \mathbf{F}_k$. As we saw above,

$$J = \langle H, q \rangle$$

and clearly $q \notin H$.

We claim that any other complementary generator of J over H is in same (H, H)-double-coset of q or of q^{-1} in J. Namely, if $J = \langle H, q' \rangle$ then $q' \in HqH \cup Hq^{-1}H$. To see this, let Y be some basis of H and think of J as the free group over the basis $Y \cup \{q\}$ (this is true because $rk(J) = rk(H) + 1$). Now think of q' as a word in the elements of this basis. Multiplying from the right or left by elements of Y does not affect the (H, H)-double-coset, so assume w.l.o.g. that q' begins and ends with either q or q^{-1}. But then the set $Y \cup \{q'\}$ is Nielsen-reduced with respect to the basis $Y \cup \{q\}$ (see, for instance, the definition in Chapter 1 of [LS70]). As consequence, $Y \cup \{q'\}$ equals $Y \cup \{q\}$ up to taking inverses (Proposition 2.8 therein). Thus $q' = q$ or $q' = q^{-1}$.

So let $\{a, b\}$ be another pair of vertices generating $\Delta(C)$. We show that it belongs to C. Let p_a, p_b be words in $\mathbf{F}_k$ corresponding to paths in $\hat{\Gamma}$ from the basepoint to

a, b respectively. We have $\langle H, p_a p_b^{-1} \rangle = J$, so $p_a p_b^{-1} \in HqH \cup Hq^{-1}H$. W.l.o.g. it is in HqH (otherwise switch a and b). So assume $p_a p_b^{-1} = h_1 q h_2$ with $h_1, h_2 \in H$. But $h_1^{-1} p_a$ is also a path from the basepoint of $\hat{\Gamma}$ to a, and likewise $h_2 p_b$ a path to b. Choosing if needed these paths instead of p_a, p_b we can assume that

$$p_a p_b^{-1} = q = p_x p_y^{-1}.$$

Thus,

$$p_a^{-1} p_x = p_b^{-1} p_y.$$

This shows that there is a path in $\hat{\Gamma}$ from a to x corresponding to a path from b to y. This shows precisely that the pair $\{a, b\}$ is in the same component of Υ as $\{x, y\}$, namely, in C. □

This completes the proof of Lemma 6.3. This Lemma, together with Lemmas 6.2 and 6.1, yields Proposition 1.9 and thus Theorem 1.5.

6.2 Further Relations between $\pi(\cdot)$ and $\phi(\cdot)$

Let us take another look now at Conjecture 1.10. It posits that the results described in Proposition 1.9 hold for all values of $\pi(\cdot)$ and $\phi(\cdot)$. To understand what this means, suppose that H is a free factor in all the quotients in $\mathcal{O}_X(H)$ of ranks up to $i-1$. What can be said about rank-i quotients in which H is a free factor? The conjecture states that their number exactly offsets the sum of two terms: The contribution to $a_i(H)$ of the quotients of smaller rank and of the term $\frac{-1}{n^{rk(H)}}$ when $i = rk(H)$. For instance, $\pi(H) = 3$ for $H = \langle x_1^2 x_2^2 x_3^2 \rangle$. In particular, H is a free factor of all quotients in $\mathcal{O}_X(H)$ of rank ≤ 2. There is a single H-critical subgroup ($\mathbf{F}_3$ itself), and additional 13 quotients of rank 3 in which H is a free factor. The contribution of quotients of rank ≤ 2 to $a_3(H)$ is indeed exactly (-13).

Interestingly enough, this is indeed the case for every free factor $H \overset{*}{\leq} \mathbf{F}_k$. In this case, since free factors are measure preserving, we get that $\phi(H) = \infty$, so $a_i(H) = 0$ for every i, and the statement of the previous paragraph holds. For the general case the conjecture states that as long as we consider low-rank quotients and "imprimitivity has not been revealed yet", the situation does not differ from what is seen in the primitive case.

We finish this section by stating another result connecting $\pi(\cdot)$ and $\phi(\cdot)$. It shows an elegant property of both of them that lends further support to our belief in Conjecture 1.10.

Two words $w_1, w_2 \in \mathbf{F}_k$ are called *disjoint* (with respect to a given basis) if they share no common letters.

Lemma 6.8. *Let $w_1, w_2 \in \mathbf{F}_k$ be disjoint. Then*

$$\pi(w_1 w_2) = \pi(w_1) + \pi(w_2)$$
$$\phi(w_1 w_2) = \phi(w_1) + \phi(w_2)$$

Moreover, $a_{\phi(w_1w_2)}(w_1w_2) = a_{\phi(w_1)}(w_1) \cdot a_{\phi(w_2)}(w_2)$*, and if part 2 of Conjecture 1.10 holds for* $H = \langle w_1 \rangle$ *and for* $H = \langle w_2 \rangle$*, then it also holds for* $H = \langle w_1 w_2 \rangle$*.*

This lemma is essentially outside the scope of the present paper, so we only sketch its proof. Let $\alpha_n \in Hom(\mathbf{F}_k, S_n)$ be a random homomorphism chosen with uniform distribution. As w_1 and w_2 are disjoint, the random permutations $\alpha_n(w_1)$ and $\alpha_n(w_2)$ are independent. The claims about the additivity of $\phi(\cdot)$ and the multiplicativity of $a_{\phi(\cdot)}(\cdot)$ are easy to derive by calculating the probability that 1 is a fixed point of w_1w_2. The key fact in this calculation is the aforementioned independence of $\alpha_n(w_1)$ and $\alpha_n(w_2)$.

The other claims in the lemma follow from an analysis of H-critical subgroups. By considering properties of the associated core graphs it is not hard to show that $J \leq \mathbf{F}_k$ is $\langle w_1w_2 \rangle$-critical iff it is the free product of a $\langle w_1 \rangle$-critical subgroup and a $\langle w_2 \rangle$-critical subgroup.

7 Primitive Words and the Profinite Completion

Most of the standard facts below about profinite groups and particularly free profinite groups can be found with proofs in [Wil98] (in particular Section 5.1).

A profinite group is a topological group G with any of the following equivalent properties:

- G is the inverse limit of an inverse system of finite groups.
- G is compact, Hausdorff and totally disconnected.
- G is isomorphic (as a topological group) to a closed subgroup of a Cartesian product of finite groups.
- G is compact and $\bigcap(N | N \lhd_O G) = 1$

The *free profinite group* on a finite set X is a profinite group F together with a map $j : X \to F$ with the following universal property: whenever $\xi : X \to G$ is a map to a profinite group G, there is a unique (continuous) homomorphism $\bar{\xi} : F \to G$ such that $\xi = \bar{\xi} j$. Such F exists for every X and is unique up to a (continuous) isomorphism. We call $j(X)$ a basis of F. It turns out that every two bases of F have the same size which is called the *rank* of F. The free profinite group of rank k is denoted by $\hat{\mathbf{F}}_k$. An element $w \in \hat{\mathbf{F}}_k$ is *primitive* if it belongs to some basis.

It is a standard fact that $\hat{\mathbf{F}}_k$ is the profinite completion of $\mathbf{F}_k$ and $\mathbf{F}_k$ is naturally embedded in $\hat{\mathbf{F}}_k$. Moreover, every basis of $\mathbf{F}_k$ is then also a basis for $\hat{\mathbf{F}}_k$, so a primitive word $w \in \mathbf{F}_k$ is also primitive as an element of $\hat{\mathbf{F}}_k$. It is conjectured that the converse also holds:

Conjecture 7.1. *A word* $w \in \mathbf{F}_k$ *is primitive in* $\hat{\mathbf{F}}_k$ *iff it is primitive in* $\mathbf{F}_k$*.*

This conjecture, if true, immediately implies the following one:

Conjecture 7.2. *The set of primitive elements in $\mathbf{F}_k$ form a closed set in the profinite topology.*

Conjecture 1.4 implies these last two conjectures (it is in fact equivalent to Conjecture 7.1, see below): we define measure preserving elements in $\hat{\mathbf{F}}_k$ as before. Namely, an element $w \in \hat{\mathbf{F}}_k$ is measure preserving if for any finite group G and a uniformly distributed random (continuous) homomorphism $\hat{\alpha}_G \in Hom(\hat{\mathbf{F}}_k, G)$, the image $\hat{\alpha}_G(w)$ is uniformly distributed in G. Clearly, an element of $\mathbf{F}_k$ is measure preserving w.r.t $\mathbf{F}_k$ iff this holds w.r.t. $\hat{\mathbf{F}}_k$.

As in the abstract case, a primitive element of $\hat{\mathbf{F}}_k$ is measure preserving. Conjecture 1.4 would therefore imply that if $w \in \mathbf{F}_k$ is primitive in $\hat{\mathbf{F}}_k$, then w is also primitive w.r.t. $\mathbf{F}_k$. In particular, Theorem 1.5 yields:

Corollary 7.3. *Let $S \subset \mathbf{F}_k$ be a finite subset of cardinality $|S| \geq k-1$. Then,*

$$S \text{ can be extended to a basis in } \hat{\mathbf{F}}_k \iff S \text{ can be extended to a basis in } \mathbf{F}_k$$

In particular, for every $w \in \mathbf{F}_2$:

$$w \text{ is primitive in } \hat{\mathbf{F}}_2 \iff w \text{ is primitive in } \mathbf{F}_2$$

This corollary yields, in turn, Corollary 1.6, which states the special case of Conjecture 7.2 for $\mathbf{F}_2$.

As shown by Chen Meiri (unpublished), Conjectures 7.1 and 1.4 are equivalent. With his kind permission we explain this result in this section. Meiri showed that in $\hat{\mathbf{F}}_k$ primitivity and measure preservation are equivalent (Proposition 7.4 below). Thus, $w \in \mathbf{F}_k$ is primitive as an element of $\hat{\mathbf{F}}_k$ iff it is measure preserving.

Proposition 7.4. *[C. Meiri, unpublished] Let w belong to $\hat{\mathbf{F}}_k$. Then*

$$w \text{ is primitive } \iff w \text{ is measure preserving}$$

Proof. The ($\Rightarrow$) implication is trivial as in the abstract case: for every finite group G and every basis $x_1, \ldots, x_k$ of $\hat{\mathbf{F}}_k$ there is a bijection

$$\begin{aligned} Hom(\hat{\mathbf{F}}_k, G) &\xrightarrow{\cong} G^k \\ \alpha_G &\mapsto (\alpha_G(x_1), \ldots, \alpha_G(x_k)) \end{aligned}$$

For the other direction, for every $w \in \hat{\mathbf{F}}_k$, finite group G and $g \in G$ define

$$\begin{aligned} H_w(G,g) &= \left\{ \alpha_G \in Hom(\hat{\mathbf{F}}_k, G) \mid \alpha_G(w) = g \right\} \\ E_w(G,g) &= \left\{ \alpha_G \in Epi(\hat{\mathbf{F}}_k, G) \mid \alpha_G(w) = g \right\} \end{aligned}$$

Now assume $w \in \hat{\mathbf{F}}_k$ is measure preserving, and let $x \in \hat{\mathbf{F}}_k$ be any primitive element. For every finite group G we have $|H_w(G,g)| = |G|^{k-1} = |H_x(G,g)|$. The same equality holds for the set of epimorphisms, namely $|E_w(G,g)| = |E_x(G,g)|$. We will show this by induction on $|G|$.

If $|G| = 1$ the claim is trivial. The inductive step goes as follows: if $g \in G$, then

$$\begin{aligned} |E_w(G,g)| &= |H_w(G,g)| - \sum_{g \in H \lneq G} |E_w(H,g)| = \\ &= |H_x(G,g)| - \sum_{g \in H \lneq G} |E_x(H,g)| = |E_x(G,g)| \end{aligned}$$

Now choose a basis $x_1, \ldots, x_k$ of $\hat{\mathbf{F}}_k$. For every $N \triangleleft_O \hat{\mathbf{F}}_k$, $|E_{x_1}(\hat{\mathbf{F}}_k/N, wN)| = |E_w(\hat{\mathbf{F}}_k/N, wN)| \geq 1$. If $\alpha \in E_{x_1}(\hat{\mathbf{F}}_k/N, wN)$ then $wN = \alpha(x_1), \alpha(x_2), \ldots, \alpha(x_k)$ generate $\hat{\mathbf{F}}_k/N$. A standard compactness argument shows that there are elements $w_2, \ldots, w_k \in \hat{\mathbf{F}}_k$ such that $\{wN, w_2N, \ldots, w_kN\}$ generate $\hat{\mathbf{F}}_k/N$ for every $N \triangleleft_O \hat{\mathbf{F}}_k$. But then $\{w, w_2, \ldots, w_k\}$ generate $\hat{\mathbf{F}}_k$ as well. Whenever k elements generate $\hat{\mathbf{F}}_k$, they generate it freely. Thus $\{w, w_2, \ldots, w_k\}$ is a basis and w is primitive. □

8 The Average Number of Fixed Points in $\alpha_n(w)$

As before, let $\alpha_n \in Hom(\mathbf{F}_k, S_n)$ be a uniformly distributed random homomorphism. In (1.1) we defined the function $\Phi_{\langle w \rangle}(n) = \Phi_w(n)$ for every $w \in \mathbf{F}_k$. It considers the probability that $\alpha_n(w)$ fixes the element 1 and quantifies its deviation from $\frac{1}{n}$. The choice of the element 1 is arbitrary, of course, and we get the same probability for every element in $1, \ldots, n$. Thus $n\Phi_w(n) + 1$ is the average number of fixed points of the random permutation $\alpha_n(w)$.

Corollary 4.2 states that in $\mathbf{F}_2$ there are exactly four possible primitivity ranks of words. This translates through Proposition 1.9 to four possibilities for the average number of fixed points in the permutation $\alpha_n(w)$, as summarized by Table 1:

$\pi(w)/\phi(w)$	Description	$Prob[\alpha_n(w)(1) = 1]$	Avg # of f.p. of $\alpha_n(w)$
0	$w = 1$	1	n
1	w is a power	$\frac{1}{n} + \frac{a_1(w)}{n} + \sum_{i=2}^{\infty} \frac{a_i(w)}{n^i}$	$1 + a_1(w) + O\left(\frac{1}{n}\right)$
2		$\frac{1}{n} + \frac{a_2(w)}{n^2} + \sum_{i=3}^{\infty} \frac{a_i(w)}{n^i}$	$1 + \frac{a_2(w)}{n} + O\left(\frac{1}{n^2}\right)$
∞	w is primitive	$\frac{1}{n}$	1

Table 1: The possibilities for the average number of fixed points of the permutation $\alpha_n(w)$ for some $w \in \mathbf{F}_2$.

Recall that all coefficients $a_i(w)$ are integers (Claim 5.1). Moreover, in these cases $a_{\phi(w)}(w)$ counts the $\langle w \rangle$-critical subgroups of $\mathbf{F}_2$, so in particular $a_{\phi(w)}(w) > 0$. We thus obtain

Corollary 8.1. *For every word $w \in \mathbf{F}_2$ and every large enough n, the average number of fixed points of $\alpha_n(w)$ is at least 1.*

This leads to the following conjecture, which is a consequence of Conjecture 1.10:

Conjecture 8.2. *For every word $w \in \mathbf{F}_k$ and every large enough n, the average number of fixed points of $\alpha_n(w)$ is at least 1.*

Proposition 1.9 says something about free words in general. If $\phi(w) \leq 2$ for some $w \in \mathbf{F}_k$, then the first non-vanishing coefficient $a_{\phi(w)}(w)$ is positive. Thus,

Corollary 8.3. *For every word $w \in \mathbf{F}_k$ the average number of fixed points in $\alpha_n(w)$ is at least $1 - O\left(\frac{1}{n^2}\right)$.*

It is suggestive to ask whether Conjecture 8.2 holds for *all* n. Namely, is it true that for every $w \in \mathbf{F}_k$ and every n, the average number of fixed points in $\alpha_n(w)$ is at least 1? By results of Abért ([Abe06]), this statement turns out to be incorrect.

A Note Added in Proof

Remark 8.4. After this paper was completed, we learned about the algorithm of Silva and Weil to detect free-factor subgroups in the free group [SW08]. In essence, their algorithm relies on the same phenomenon that we independently noticed here. However, our reasoning is very different, and offers several substantial advantages over the presentation in [SW08]. A more elaborate discussion of the differences between the two approaches appears in Appendix A.

Remark 8.5. In subsequent joint work with O. Parzanchevski [PP12], we manage to prove Conjecture 1.4 in full. That proof relies on Theorem 1.1 and follows the general strategy laid out in the current paper. In particular, we establish Conjectures 1.10, 7.1, 7.2 and 8.2.

Acknowledgements

It is a pleasure to thank Nati Linial for his support, encouragement and useful comments. We are also grateful to Aner Shalev for supporting this research and for his valuable suggestions. We would also like to thank Tsachik Gelander, Michael Larsen, Alex Lubotzky, Chen Meiri, Ori Parzanchevski, Iddo Samet and Enric Ventura for their beneficial comments. We would also like to express our gratefulness to the anonymous referee for his many valuable comments.

Appendices

A An Algorithm to Detect Free Factors

One of the interesting usages of Theorem 1.1 is an algorithm to detect free factor subgroups and consequently, also to detect primitive words in $\mathbf{F}_k$. The algorithm receives as input H and J, two finitely generated subgroups of $\mathbf{F}_k$, and determines whether $H \overset{*}{\leq} J$. The subgroups H and J are given to us by specifying a generating set, where members of the generating sets are words in the elements of the fixed basis X. (Note that the algorithm in particular decides as well whether $H \leq J$, but this is neither hard nor new).

We should mention that ours is not the first algorithm, nor the first graph-theoretic one, for this problem (see Chapter I.2 in [LS70]). We already mentioned (Remark 8.4) [SW08], who noticed the basic phenomenon underlying our algorithm, albeit in a very different language. See Remark A.2 below for an explanation of the differences. A well-known algorithm due to Whitehead solves a much more general problem. Namely, for given $2r$ words $w_1, \ldots, w_r, u_1, \ldots, u_r \in \mathbf{F}_k$, it determines whether there is an automorphism $\alpha \in Aut(\mathbf{F}_k)$ such that $\alpha(w_i) = u_i$ for each i ([Whi36a],[Whi36b]. For a good survey see Chapter I.4 at [LS70]. A nice presentation of the restriction of Whitehead's algorithm to our problem appears in [Sta99]). Quite recently, Roig, Ventura and Weil introduced a more clever version of the Whitehead algorithm for the case of detecting primitive words and free factor subgroups [RVW07]. Their version of the algorithm has polynomial time in both the length of the given word w (or the total length of generators of a given subgroup H) and in k, the rank of the ambient group $\mathbf{F}_k$. To the best of our knowledge, their algorithm is currently the best one for this problem, complexity-wise. The algorithm we present is, at least naively, exponential, as we show below (Remark A.1).

So assume we are given two subgroups of finite rank of $\mathbf{F}_k$, H and J, by means of finite generating sets S_H, S_J. Each element of S_H, S_J is assumed to be a word in the letters $X \cup X^{-1}$ (recall that $X = \{x_1, \ldots, x_k\}$ is the given basis of $\mathbf{F}_k$). To find out whether $H \overset{*}{\leq} J$, follow the following steps.

Step 1: Construct Core Graphs and Morphism
First, construct the core graphs $\Gamma = \Gamma_X(H)$ and $\Delta = \Gamma_X(J)$ by the process described in Appendix C. Then, seek a morphism $\eta : \Gamma \to \Delta$. This is a simple process that can be done inductively as follows: η must map the basepoint of Γ to the basepoint of Δ. Now, as long as η is not fully defined, there is some j-edge $e = (u, v)$ in $E(\Gamma)$ for which the image is not known yet, but the image of one of the end points, say $\eta(u)$, is known (recall that Γ is connected). There is at most one possible value that $\eta(e)$ can take, since the star of $\eta(u)$ contains at most one outgoing j-edge. If there is no such edge, we get stuck. Likewise, $\eta(v)$ must equal the terminus of $\eta(e)$, and

if $\eta(v)$ was already determined in an inconsistent way, we get a contradiction. If in this process we never get stuck and never reach a contradiction, then η is defined. If this process cannot be carried out, then there is no morphism from Γ to Δ, and hence H is not a subgroup of J (see Claim 2.2).

Step 2: Reduce to Two Groups with $H \overset{X}{\twoheadrightarrow} J'$
After constructing the morphism $\eta : \Gamma \to \Delta$, we obtain a new graph from Δ by omitting all edges and all vertices not in the image of η. Namely,

$$\Delta' := \eta(\Gamma)$$

It is easy to see that Δ' is a core-graph, and we denote by J' the subgroup corresponding to Δ'. Obviously, Δ' is a quotient of Γ, so $H \overset{X}{\twoheadrightarrow} J'$. Moreover, it follows from Claim 2.5 that

$$H \overset{*}{\leq} J \iff H \overset{*}{\leq} J'.$$

Step 3: Use $\rho_X(H, J')$ to determine whether $H \overset{*}{\leq} J'$
Now calculate $\rho_X(H, J')$ (this is clearly doable because the subgraph of $\mathcal{D}_k$ consisting of quotients of Γ is finite). Thanks to Theorem 1.1, $\rho_X(H, J')$ determines whether or not $H \overset{*}{\leq} J'$, and consequently, whether or not $H \overset{*}{\leq} J$.

Remark A.1. The complexity of this algorithm is roughly $O(v^{2t})$, where v is the number of vertices in $\Gamma_X(H)$ and t is the difference in ranks: $t = rk(J) - rk(H)$. Naively, we need to go over roughly all possible sets of t pairs of vertices of $\Gamma_X(H)$ and try to merge them (see Remark 3.7). The number of possibilities is at most $\binom{\binom{v}{2}}{t}$, which shows the claimed bound. (In fact, we can restrict to pairs where both vertices are in the same fiber of the morphism $\eta : \Gamma_X(H) \to \Gamma_X(J)$.)

A.1 Examples

We illustrate the different phases of the algorithm by two concrete examples. Consider first the groups $H = \langle x_1 x_2 {x_1}^{-1} {x_2}^{-1}, x_2 {x_1}^2 \rangle$ and $J = \langle {x_1}^3, {x_2}^3, x_1 {x_2}^{-1}, x_1 x_2 x_1 \rangle$, both in $\mathbf{F}_2$. The core graphs of these groups are:

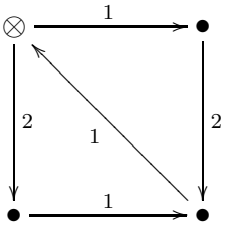

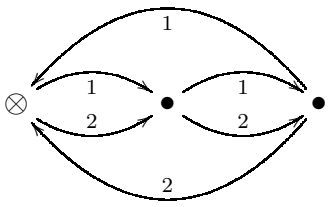

In this case, a morphism η from $\Gamma = \Gamma_X(H)$ to $\Delta = \Gamma_X(J)$ can be constructed. All the vertices of Γ are in the image of η, and only one edge, the long 2-edge at the bottom, is not in $\eta(E(\Gamma))$. Thus Δ' is:

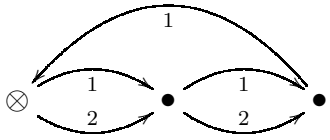

and J' is the corresponding subgroup $J' = \langle x_1^3, x_1 x_2^{-1}, x_1 x_2 x_1 \rangle$.

Finally, $rk(H) = 1 - \chi(\Gamma) = 2$ and $rk(J') = 1 - \chi(\Delta') = 3$, and so the difference is $rk(J') - rk(H) = 1$. It can be easily verified that Δ' is indeed an immediate quotient of Γ: simply merge the upper-right vertex of Γ with the bottom-left one to obtain Δ'. Thus $\rho_X(H, J') = 1 = rk(J') - rk(H)$, and so $H \overset{*}{\leq} J'$ hence $H \overset{*}{\leq} J$.

As a second example, consider the commutator word $w = x_1 x_2 x_1^{-1} x_2^{-1}$. We want to determine whether it is primitive in $\mathbf{F}_3$. We take $H = \langle w \rangle$ and the core graphs are then

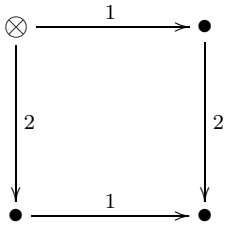

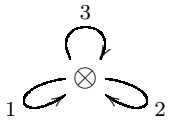

Once again, a morphism η from $\Gamma = \Gamma_X(H)$ to $\Delta = \Gamma_X(\mathbf{F}_3)$ can be constructed, and there is a single edge in Δ, the 3-edge, outside the image of η. Thus Δ' is the quotient of Γ which is the bottom graph in Figure 3.1, and J' is simply $\mathbf{F}_2$.

Finally, $rk(H) = 1 - \chi(\Gamma) = 1$ and $rk(\mathbf{F}_2) = 2$, and so the difference is $rk(F_2) - rk(H) = 1$. But as we infer from Figure 3.1, $\rho_X(H, \mathbf{F}_2) = 2$. Thus $\rho_X(H, \mathbf{F}_2) > rk(\mathbf{F}_2) - rk(H)$ and H is *not* a free factor of $\mathbf{F}_2$. As consequence, w is *not* primitive in $\mathbf{F}_3$. (This example generalizes as follows: if w is a free word containing exactly l different letters, then w is primitive iff we can obtain a wedge-of-loops graph from $\Gamma_X(\langle w \rangle)$ by merging $l - 1$ pairs of vertices.)

Remark A.2. At this point we would like to elaborate on the differences between the algorithm presented here and the one introduced in [SW08]. Silva and Weil's presentation considers automata and their languages. We consider the X-fringe $\mathcal{O}_X(H)$ and introduce the DAG $\mathcal{D}_k$ and the distance function from Definition 3.2. Steps 1 and 2 of our algorithm, which reduce the problem in its very beginning to the

case where $H \overset{X}{\twoheadrightarrow} J$, have no parallel in [SW08]. Rather, they show that if $H \overset{*}{\leq} J$, then by some sequence of "i-steps" (their parallel of our immediate quotients) on H, of length at most $rk(J) - rk(H)$, one can obtain a core graph which is embedded in $\Gamma_X(J)$ (we make the observation that this embedded core graph can be computed in advance). Besides shedding more light on this underlying phenomenon, our more graph-theoretic approach has another substantial advantage: by considering $\mathcal{D}_k$, turning the fringe $\mathcal{O}_X(H)$ into a directed graph and stating the algorithm in the language of Theorem 1.1, we obtain a straight-forward algorithm to identify H-critical subgroups and to compute $\pi(H)$. Moreover, we obtain a straight-forward algorithm to identify all "algebraic extensions" of H (Corollary 4.4). In particular, our algorithm to identify algebraic extensions substantially improves the one suggested in [KM02], Theorem 11.3 (and see also remark 11.4 about its efficiency).

B The Proof of Lemma 3.3

To complete the picture, we prove the upper bound for $\rho_X(H, J)$ stated in Lemma 3.3. We need to show that if $H, J \leq_{fg} \mathbf{F}_k$ such that $H \overset{X}{\twoheadrightarrow} J$, then

$$\rho_X(H, J) \quad \leq \quad \mathrm{rk}(J)$$

Proof. We show that $\Delta = \Gamma_X(J)$ can be obtained from $\Gamma = \Gamma_X(H)$ by merging at most $\mathrm{rk}(J)$ pairs of vertices. To see this, denote by m the number of edges in Γ, and choose some order on these edges, $e_1, \ldots, e_m$ so that for every i, there is a path from the basepoint of Γ to e_i traversing only edges among $e_1, \ldots, e_{i-1}$. (So e_1 must be incident with the basepoint, e_2 must be incident either with the basepoint or with the other end of e_1, etc.)

We now expose Δ step by step, each time adding the images of the next edge of Γ and of its end points. Formally, denote by η the (surjective) morphism from Γ to Δ, let Γ_i be the subgraph of Γ that is the union of the basepoint of Γ together with $e_1, \ldots, e_i$ and their endpoints, and let $\Delta_i = \eta(\Gamma_i)$. We thus have two series of subgraphs

$$\Gamma_0 \subseteq \Gamma_1 \subseteq \ldots \subseteq \Gamma_m = \Gamma$$

and

$$\Delta_0 \subseteq \Delta_1 \subseteq \ldots \subseteq \Delta_m = \Delta$$

with $\Delta_0 = \Gamma_0$ being graphs with a single vertex and no edges.

Assume that $e_i = (u, v)$, and w.l.o.g. that $u \in V(\Gamma_{i-1})$. We then distinguish between three options. A **forced** step is when $\eta(e_i)$ already belongs to Δ_{i-1} and then $\Delta_i = \Delta_{i-1}$. A **free** step takes place when neither $\eta(e_i)$ nor $\eta(v)$ belong to Δ_{i-1}, in which case $\pi_1(\Delta_i) = \pi_1(\Delta_{i-1})$. The third option is that of a **coincidence**. This means that $\eta(e_i)$ does not belong to Δ_{i-1} but $\eta(v)$ does. In this case, Δ_i is obtained from Δ_{i-1} by connecting two vertices by a new edge, and $\pi_1(\Delta_i)$ has rank larger by

1 from the rank of $\pi_1(\Delta_{i-1})$. Since the fundamental group of Δ_0 has rank 0, this shows there are exactly $rk(J)$ coincidences along this process.

Assume the coincidences occurred in steps $j_1, \ldots, j_{\mathrm{rk}(J)}$. If $e_{j_i} = (u, v)$, we let $\tilde{v} \in \eta^{-1}(\eta(v)) \cap V(\Gamma_{i-1})$, and take $\{v, \tilde{v}\}$ to be a pair of vertices of Γ that we merge. (It is possible that $v = w$.) Let $w_i \in \mathbf{F}_k$ be be a word corresponding to this merge in Γ. It is easy to see by induction that Δ_{j_i} corresponds to the subgroup $\langle H, w_1, \ldots, w_i \rangle$. In particular, Δ corresponds to $\left\langle H, w_1, \ldots, w_{\mathrm{rk}(\mathrm{J})} \right\rangle$. We are done because all these words correspond to pairs of vertices in Γ (and see Remark 3.7). □

C The Folding Algorithm to Construct Core Graphs

Finally, we present a well known algorithm to construct the core graph of a given subgroup $H \leq_{fg} \mathbf{F}_k$. The input to this process is any finite set of words $\{h_1, \ldots, h_r\}$ in the letters $\{x_1, \ldots, x_k\}$ that generate H.

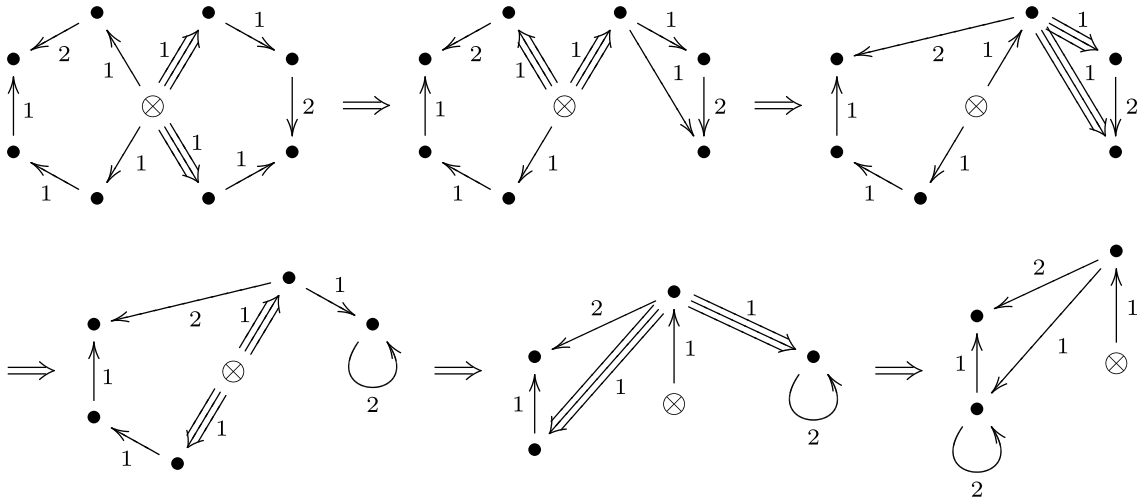

Figure C.1: Generating the core graph $\Gamma_X(H)$ of $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle \leq \mathbf{F}_2$ from the given generating set. We start with the upper left graph which contains a distinct loop at the basepoint for each (reduced) element of the generating set. Then, gradually and at arbitrary order, we merge pairs of equally-labeled edges which share the same origin or the same terminus. One of the possible orders of merging pairs is shown in this figure, and at each phase we mark by triple arrows the pair of edges being merged. The graph in the bottom right is $\Gamma_X(H)$, as it has no equally-labeled edges sharing the same origin or the same terminus.

Every element h_i of the generating set corresponds to some path with directed edges labeled by the x_i's (we assume the element is given in reduced form). Merge these r paths to a single graph by identifying all their $2r$ end-points to a single vertex

which is denoted as basepoint. Then, as long as there are two j-labeled edges with the same terminus (resp. origin) for some j, merge the two edges and their origins (resp. termini). Such a step is often referred to as a *Stallings' folding*. It is a fairly easy observation that the resulting graph is indeed $\Gamma_X(H)$ and that the order of folding has no significance. To illustrate, we draw in Figure C.1 the folding process by which we obtain the core graph $\Gamma_X(H)$ of $H = \langle x_1 x_2 x_1^{-3}, x_1^2 x_2 x_1^{-2} \rangle \leq \mathbf{F}_2$ from the given generating set.

References

[Abe06] Miklós Abert, *On the probability of satisfying a word in a group*, Journal of Group Theory **9** (2006), 685–694. 8

[Bog08] Oleg Bogopolski, *Introduction to group theory*, EMS Textbooks in Mathematics, European Mathematical Society, Zurich, 2008. 3.1

[Ger84] SM Gersten, *On whitehead's algorithm*, Bull. Amer. Math. Soc., New Ser **10** (1984), no. 2, 281–284. 1.2

[GS09] Shelly Garion and Aner Shalev, *Commutator maps, measure preservation, and t-systems*, Trans. Amer. Math. Soc. **361** (2009), no. 9, 4631–4651. 1

[KM02] I. Kapovich and A. Myasnikov, *Stallings foldings and subgroups of free groups*, Journal of Algebra **248** (2002), no. 2, 608–668. 1, 2.1, 2.2, 4, A.2

[LP10] Nati Linial and Doron Puder, *Words maps and spectra of random graph lifts*, Random Structures and Algorithms **37** (2010), no. 1, 100–135. 1, 1, 1, 1, 5, 6.1

[LS70] R.C Lyndon and P.E. Schupp, *Combinatorial group theory*, Springer-Verlag, Berlin; New York, 1970. 6.1, A

[LS08] Michael Larsen and Aner Shalev, *Characters of symmetric groups: sharp bounds and applications*, Inventiones mathematicae **174** (2008), no. 3, 645–687. 1

[LS09] ———, *Words maps and waring type problems*, J. Amer. Math. Soc. **22** (2009), no. 2, 437–466. 1

[MVW07] Alexei Miasnikov, Enric Ventura, and Pascal Weil, *Algebraic extensions in free groups*, Geometric group theory (G.N. Arzhantseva, L. Bartholdi, J. Burillo, and E. Ventura, eds.), Trends Math., Birkhauser, 2007, pp. 225–253. 1, 1, 1, 2.1, 2.2, 2.3, 4

[Nic94] Alexandru Nica, *On the number of cycles of given length of a free word in several random permutations*, Random Structures and Algorithms **5** (1994), no. 5, 703–730. 1, 1, 1, 5

[PP12] D. Puder and O. Parzanchevski, *Measure preserving words are primitive*, Arxiv preprint arXiv:1202.3269 (2012). 8.5

[RVW07] A. Roig, E. Ventura, and P. Weil, *On the complexity of the whitehead minimization problem*, International journal of Algebra and Computation **17** (2007), no. 8, 1611–1634. A

[Seg09] Dan Segal, *Words: notes on verbal width in groups*, London Mathematical Society, Lecture note Series 361, Cambridge University Press, Cambridge, 2009. 1, 1

[Sha09] Aner Shalev, *Words maps, conjugacy classes, and a non-commutative waring-type theorem*, Annals of Math. **170** (2009), 1383–1416. 1

[Sta83] John R. Stallings, *Topology of finite graphs*, Inventiones mathematicae **71** (1983), no. 3, 551–565. 1, 2.1, 2.2

[Sta99] ———, *Whitehead graphs on handlebodies*, Geometric group theory down under (J. Cossey, C. F. Miller, W.D. Neumann, and M. Shapiro, eds.), de Gruyter, Berlin, 1999, pp. 317–330. 1.2, A

[SW08] P. Silva and P. Weil, *On an algorithm to decide whether a free group is a free factor of another*, RAIRO - Theoretical Informatics and Applications **42** (2008), no. 2, 395–414. 8.4, A, A.2

[Whi36a] J.H.C. Whitehead, *On certain sets of elements in a free group*, Proc. London Math. Soc. **41** (1936), 48–56. 1.2, A

[Whi36b] ———, *On equivalent sets of elements in a free group*, Ann. of Math. **37** (1936), 768–800. 1.2, A

[Wil98] John S. Wilson, *Profinite groups*, Clarendon Press, Oxford, 1998. 7